

## Power ISA ${ }^{\text {TM }}$ Version 3.0

November 30, 2015

## IBM

© Copyright International Business Machines Corporation 1994-2015. All rights reserved.

Printed in the United States of America November, 2015

By downloading the POWER® Instruction set Architecture ("ISA") Specification, you agree to be bound by the terms and conditions of this agreement.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Other company, product, and service names may be trademarks or service marks of others.

All information contained in this document is subject to change without notice. The products described in this document are NOT intended for use in applications such as implantation, life support, or other hazardous uses where malfunction could result in death, bodily injury, or catastrophic property damage. The information contained in this document does not affect or change IBM product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of IBM or third parties. All information contained in this document was obtained in specific environments, and is presented as an illustration. The results obtained in other operating environments may vary.

While the information contained herein is believed to be accurate, such information is preliminary, and should not be relied upon for accuracy or completeness, and no representations or warranties of accuracy or completeness are made.

Note: This document contains information on products in the design, sampling and/or initial production phases of development. This information is subject to change without notice. Verify with your IBM field applications engineer that you have the latest version of this document before finalizing a design.

You may use this documentation solely for developing technology products compatible with Power Architecture ${ }^{\circledR}$ in support of growing the POWER ecosystem. You may not modify this documentation. You may distribute the documentation to suppliers and other contractors hired by you solely to produce your technology products compatible with Power Architecture $®$ technology and to your customers (either directly or indirectly through your resellers) in conjunction with their use and instruction of your technology products compatible with Power Architecture ${ }^{\circledR}$ technology. This agreement does not include rights to create a CPU design to run
the POWER ISA unless such rights have been granted by IBM under a separate agreement. The POWER ISA specification is protected by copyright and the practice or implementation of the information herein may be protected by one or more patents or pending patent applications. No other license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document.

THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED ON AN "AS IS" BASIS. IBM makes no representations or warranties, either express or implied, including but not limited to, warranties of merchantability, fitness for a particular purpose, or non-infringement, or that any practice or implementation of the IBM documentation will not infringe any third party patents, copyrights, trade secrets, or other rights. In no event will IBM be liable for damages arising directly or indirectly from any use of the information contained in this document.

IBM Systems and Technology Group
2070 Route 52, Bldg. 330
Hopewell Junction, NY 12533-6351
The IBM home page can be found at ibm.com ${ }^{\circledR}$.

The following paragraph does not apply to the United Kingdom or any country or state where such provisions are inconsistent with local law.

The specifications in this manual are subject to change without notice. This manual is provided "AS IS". International Business Machines Corp. makes no warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

International Business Machines Corp. does not warrant that the contents of this publication or the accompanying source code examples, whether individually or as one or more groups, will meet your requirements or that the publication or the accompanying source code examples are error-free.

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication.

Address comments to IBM Corporation, 11400 Burnett Road, Austin, Texas 78758-3493. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.
The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

```
IBM^
Power ISA
PowerPC®
Power Architecture
PowerPC Architecture
Power Family
RISC/System 6000^^
POWER®
POWER2
POWER4
POWER4+
POWER5
POWER5+
POWER6®
POWER7®
System/370
System z
```

The POWER ARCHITECTURE and POWER.ORG. word marks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
AltiVec is a trademark of Freescale Semiconductor, Inc. used under license.

Notice to U.S. Government Users-Documentation Related to Restricted Rights-Use, duplication or disclosure is subject to restrictions set fourth in GSA ADP Schedule Contract with IBM Corporation.

## Preface

The roots of the Power ISA (Instruction Set Architecture) extend back over a quarter of a century, to IBM Research. The POWER (Performance Optimization With Enhanced RISC) Architecture was introduced with the RISC System/6000 product family in early 1990. In 1991, Apple, IBM, and Motorola began the collaboration to evolve to the PowerPC Architecture, expanding the architecture's applicability. In 1997, Motorola and IBM began another collaboration, focused on optimizing PowerPC for embedded systems, which produced Book E.

In 2006, Freescale and IBM collaborated on the creation of the Power ISA Version 2.03, which represented the reunification of the architecture by combining Book E content with the more general purpose PowerPC Version 2.02.
Power ISA Version 3.0 continues this integration by providing a single Book III for all Power implementations, and also by eliminating optional architecture categories to ensure increased application portability between Power processors.

The Power ISA Version 3.0 is comprised of three books and a set of appendices.
Book I, Power ISA User Instruction Set Architecture, covers the base instruction set and related facilities available to the application programmer. It includes five chapters derived from APU function, including the vector extension also known as Altivec.

Book II, Power ISA Virtual Environment Architecture, defines the storage model and related instructions and facilities available to the application programmer.
Book III, Power ISA Operating Environment Architecture, defines the supervisor instructions and related facilities used for general purpose implementations.
As used in this document, the term "Power ISA" refers to the instructions and facilities described in Books I, II, and III.

Change bars have been included to indicate changes from the Power ISA Version 2.07B, except that change bars may be omitted for changes associated with removing obsolete material.

## Summary of Changes in Power ISA Version 3.0

This document is Version 3.0 of the Power ISA. It is intended to supersede and replace version 2.07B. Any product descriptions that reference a version of the architecture are understood to reference the latest version. This version was created by making miscellaneous corrections and by applying the following requests for change (RFCs) to Power ISA Version 2.07B.

Instruction Fusion: Specifies instruction sequences that, when placed consecutively in the program, are expected to provide improved performance.
Hashing Support Operations: Adds new Count Trailing Zeros and Modulo instructions

Decimal Integer Support Operations: Adds new BCD support instructions, including variable-length load/ store instructions for bcd values, new format conversion instructions between BCD and National decimal, zoned decimal, and 128 -bit signed integer formats. new BCDtruncate, round, and shift instructions, new BCD sign digit manipulation instructions. Also adds multi-ply-by-10 instructions to faciliate binary-to-decimal conversion for printf.
Decimal Floating-Point Support Operations: Add immediate forms of DFP Test Significance instructions.

Binary Floating-Point Support Operations: Adds new binary floating-point support instructions (e.g., exponent and significand extraction and insertion) to enhance implementation of math libraries.

Quad-Precision Binary Floating-Point Operations: Add new instructions to support IEEE-754-2008 binary128 floating-point.

String Operations (FXU option): Adds instructions to accelerate character testing functions.
String Operations (VSU option): Adds instructions to accelerate string processing and targeted character extraction.

Vector Half-Precision Floating-Point Support Operations: Adds support for IEEE-754-2008 binary16 float-ing-point as a transport format.
128-bit SIMD Video Compression Operations: Adds instructions to accelerate motion estimation.

128-bit SIMD FXU Operations: Adds remaining 32-bit and 64-bit FXU functionality to vector instruction set.
128-bit SIMD Miscellaneous Operations: Enhances support for Little-Endian processing with new load/store instructions and new permute-class instructions, new byte and halfword element load/store instructions, and vector element insertion/extraction.

System Call Extension: Provides a new form of system call that can direct execution to one of a number of locations and that provides other enhancements.

PC-Relative Addressing: Specifies a new instruction that adds an immediate value to the program counter and writes it to the destination register in preparation for use with a D-Form Load instructon.

Hypervisor msgsnd Instruction Enhancements: Extends the msgsnd instruction so that messages can be sent throughout the system.

Performance Monitor Enhancements: Adds additional BHRB filtering modes, reserves a special no-op instruction for use by the Performance Monitor, and increases the scope of control of the Performance Monitor bit of the Hypervisor Facility Status and Control register.

Radix Tree and Related MMU Extensions: Adds support for the radix tree style of MMU with full virtualization and related control mechanisms that manage its coexistence with the HPT. Also adds atomic hardware refererence and change bit updates and a tlbie variant that invalidates multiple consecutive translations or translations in a range of address space.

Copy-Paste Facility: Adds support for a new facility that enables an application to move a line of data to or from memory owned by another program, possibly in another partition or system. A variant is used to initiate accelerator operations.

Optimizing mtspr Sequences: Reserves an SPR to be used in a noop mtspr to indicate the beginning of a sequence of mtsprs that can be done without synchronizing each one independently.

Atomic Memory Operations: Adds support for a new facility that performs simple atomic operations directly in memory to avoid bringing the line through the cache hierarchy when another core is likely to be the next user.

Event-Based Branch Extension: Adds External Event-Based Branch exception and status bits to the BESCR.

Processor Compatibility Register: Adds a new V 2.07 bit to the PCR that controls the availability facilities in problem state that are introduced in this level of the architecture.

Atomicity and Alignment Enhancements: Limits the number of disjoint atomic storage accesses that are allowed for various non-atomic storage accesses.
Load Doubleword Monitored Instruction: Specifies a new instruction and facility that improves the performance of the garbage collection process and eliminates the need to pause all applications during collection.

Power-Saving Mode: Replaces the existing power-saving mode instructions with a single stop instruction, and enables the operating system to enter a limited set of power-saving levels without hypervisor involvement.

D-form VSX Floating-Point Storage Access Instructions: Adds base+displacement forms of VSR load and store instructions.

Integer Multiply-Add Instructions: Adds new integer multiply-add instructions to accelerate arbitrary-length multiplication.
msgsndp Hypervisor Facility Availability Interrupt: Adds a new HFSCR bit to control the availability of the $\boldsymbol{m s g s n d p}$ instruction and the associated control registers.

VSX Permute: Adds new pernute instructions that can address all 64 VSRs.

Array Index Support: Enhance support for mixed-datatype addressing into arrays (e.g., base + 32-bit index)

Hypervisor Virtualization Interrupt: Defines a new exception and corresponding interrupt that is caused by events external to the processor that relate to virtualization.

Debug Extension for TM: Adds a new type of Trace interrupt to enable skipping past a transaction in the debugging process.
wait Instruction Enhancements: Improves the capabilities of the wait instruction so that resumption of processing can occur due to event-based branches and external signals.

Decrementer and Hypervisor Decrementer Enahncements: Defines a new mode bit in the LPCR that enables additional Decrementer and Hypervisor Decrementer bits in order to increase the time between the associated interrupts.
Deliver A Random Number: Adds a new instruction to place a random number in a GPR in one of three formats.

Data Storage Interrupt Status Register for Alignment Interrupt: Simplifies the Alignment interrupt by removing the Data Storage Interrupt Status Register (DSISR) from the set of registers modified by the Alignment interrupt.

CA32 \& OV32 and Move XER to CR Extended: Added support for 32-bit CA \& OV status in 64-bit mode for dynamically-typed languages

VSX Shift Variable: Accelerate parallel element extraction from packed vectors of arbitrary-width-element values.

Enhanced Virtualization for Linux: Delivers exceptions caused by the OS attempting to use hypervisor instructions and SPRs to the hypervisor instead of the OS. Accesses to unimplemented SPRs by the OS newly cause interrupts that are also directed to the hypervisor.
Synchronizing Messages and Storage Updates: Adds a new instruction to make latent storage updates from another thread accessible after receiving a Directed Hypervisor Doorbell interrupt from that thread.

Version 3.0

## Table of Contents

Preface ..... v
Summary of Changes in Power ISA Ver-
Table of Contents ..... ix
Book I:
Power ISA User Instruction Set Architecture ..... 1
Chapter 1. Introduction ..... 3
1.1 Overview ..... 3
1.2 Instruction Mnemonics and Operands3
1.3 Document Conventions ..... 3
1.3.1 Definitions ..... 3
1.3.2 Notation ..... 4
1.3.3 Reserved Fields, Reserved Values, and Reserved SPRs ..... 5
1.3.4 Description of Instruction Operation6
1.3.5 Phased-Out Facilities. ..... 8
1.4 Processor Overview ..... 9
1.5 Computation modes ..... 10
1.6 Instruction Formats ..... 11
1.6.1 A-FORM ..... 12
1.6.2 B-FORM ..... 12
1.6.3 D-FORM ..... 12
1.6.4 DQ-FORM ..... 12
1.6.5 DS-FORM ..... 12
1.6.6 DX-FORM ..... 12
1.6.7 I-FORM ..... 12
1.6.8 M-FORM ..... 12
1.6.9 MD-FORM ..... 12
1.6.10 MDS-FORM ..... 12
1.6.11 SC-FORM ..... 12
1.6.12 VA-FORM ..... 12
1.6.13 VC-FORM ..... 12
1.6.14 VX-FORM ..... 13
1.6.15 X-FORM ..... 13
1.6.16 XFL-FORM ..... 15
1.6.17 XFX-FORM ..... 15
1.6.18 XL-FORM ..... 15
1.6.19 XO-FORM ..... 15
1.6.20 XS-FORM ..... 15
1.6.21 XX2-FORM ..... 15
1.6.22 XX3-FORM ..... 15
1.6.23 XX4-FORM ..... 15
1.6.24 Z22-FORM ..... 15
1.6.25 Z23-FORM ..... 16
1.7 Instruction Fields ..... 16
1.8 Classes of Instructions ..... 22
1.8.1 Defined Instruction Class ..... 22
1.8.2 Illegal Instruction Class ..... 22
1.8.3 Reserved Instruction Class ..... 22
1.9 Forms of Defined Instructions ..... 23
1.9.1 Preferred Instruction Forms ..... 23
1.9.2 Invalid Instruction Forms ..... 23
1.9.3 Reserved-no-op Instructions ..... 23
1.10 Exceptions. ..... 23
1.11 Storage Addressing ..... 24
1.11.1 Storage Operands ..... 24
1.11.2 Instruction Fetches ..... 26
1.11.3 Effective Address Calculation ..... 27
Chapter 2. Branch Facility ..... 29
2.1 Branch Facility Overview ..... 29
2.2 Instruction Execution Order. ..... 29
2.3 Branch Facility Registers. ..... 30
2.3.1 Condition Register ..... 30
2.3.2 Link Register ..... 32
2.3.3 Count Register ..... 32
2.3.4 Target Address Register. ..... 32
2.4 Branch History Rolling Buffer ..... 32
2.4.1 Branch History Rolling Buffer Entry Format ..... 33
2.5 Branch Instructions ..... 34
2.6 Condition Register Instructions ..... 41
2.6.1 Condition Register Logical Instruc- tions. ..... 41
2.6.2 Condition Register Field Instruction. ..... 42
2.7 System Call Instructions. ..... 43
2.8 Branch History Rolling Buffer Instruc- tions. ..... 44
Chapter 3. Fixed-Point Facility ..... 45
3.1 Fixed-Point Facility Overview ..... 45
3.2 Fixed-Point Facility Registers ..... 45
3.2.1 General Purpose Registers ..... 45
3.2.2 Fixed-Point Exception Register ..... 45
3.2.3 VR Save Register ..... 46
3.2.4 Load Monitored Region Registers. 46
3.2.4.1 Load Monitored Region Register46
3.2.4.2 Load Monitored Section Enable
Register ..... 46
3.3 Fixed-Point Facility Instructions ..... 47
3.3.1 Fixed-Point Storage Access Instruc- tions ..... 47
3.3.1.1 Storage Access Exceptions ..... 47
3.3.2 Fixed-Point Load Instructions ..... 47
3.3.2.1 64-bit Fixed-Point Load Instruc- tions ..... 52
3.3.3 Fixed-Point Store Instructions ..... 55
3.3.3.1 64-bit Fixed-Point Store Instruc- tions ..... 58
3.3.4 Fixed Point Load and Store Quad- word Instructions ..... 59
3.3.5 Fixed-Point Load and Store with ByteReversal Instructions61
3.3.5.1 64-Bit Load and Store with Byte Reversal Instructions ..... 62
3.3.6 Fixed-Point Load and Store Multiple Instructions ..... 63
3.3.7 Fixed-Point Move Assist Instructions [Phased Out]. ..... 64
3.3.8 Other Fixed-Point Instructions ..... 67
3.3.9 Fixed-Point Arithmetic Instructions 68
3.3.9.1 64-bit Fixed-Point Arithmetic
Instructions ..... 80
3.3.10 Fixed-Point Compare Instructions.85
3.3.10.1 Character-Type Compare Instruc-
tions ..... 87
3.3.11 Fixed-Point Trap Instructions ..... 89
3.3.11.1 64-bit Fixed-Point Trap Instruc- tions ..... 90
3.3.12 Fixed-Point Select ..... 90
3.3.13 Fixed-Point Logical Instructions ..... 91
3.3.13.1 64-bit Fixed-Point Logical Instruc- tions ..... 98
3.3.14 Fixed-Point Rotate and Shift Instructions ..... 100
3.3.14.1 Fixed-Point Rotate Instructions100
3.3.14.1.1 64-bit Fixed-Point RotateInstructions103
3.3.14.2 Fixed-Point Shift Instructions ..... 106
3.3.14.2.1 64-bit Fixed-Point Shift Instruc-
tions ..... 108
3.3.15 Binary Coded Decimal (BCD)Assist Instructions ................... 110
3.3.16 Move To/From Vector-Scalar Regis- ter Instructions ..... 111
3.3.17 Move To/From System Register Instructions ..... 116
Chapter 4. Floating-Point Facility 123
4.1 Floating-Point Facility Overview ..... 123
4.2 Floating-Point Facility Registers ..... 124
4.2.1 Floating-Point Registers ..... 124
4.2.2 Floating-Point Status and Control Register. ..... 124
4.3 Floating-Point Data ..... 127
4.3.1 Data Format ..... 127
4.3.2 Value Representation ..... 127
4.3.3 Sign of Result ..... 129
4.3.4 Normalization and Denormalization ..... 129
4.3.5 Data Handling and Precision ..... 129
4.3.5.1 Single-Precision Operands ..... 129
4.3.5.2 Integer-Valued Operands ..... 130
4.3.6 Rounding ..... 131
4.4 Floating-Point Exceptions ..... 132
4.4.1 Invalid Operation Exception ..... 134
4.4.1.1 Definition ..... 134
4.4.1.2 Action ..... 134
4.4.2 Zero Divide Exception ..... 134
4.4.2.1 Definition ..... 134
4.4.2.2 Action ..... 135
4.4.3 Overflow Exception ..... 135
4.4.3.1 Definition ..... 135
4.4.3.2 Action ..... 135
4.4.4 Underflow Exception ..... 136
4.4.4.1 Definition ..... 136
4.4.4.2 Action ..... 136
4.4.5 Inexact Exception ..... 136
4.4.5.1 Definition ..... 136
4.4.5.2 Action ..... 136
4.5 Floating-Point Execution Models ..... 137
4.5.1 Execution Model for IEEE Opera- tions ..... 137
4.5.2 Execution Model for Multiply-Add Type Instructions ..... 139
4.6 Floating-Point Facility Instructions ..... 140
4.6.1 Floating-Point Storage Access Instructions ..... 141
4.6.1.1 Storage Access Exceptions ..... 141
4.6.2 Floating-Point Load Instructions ..... 141
4.6.3 Floating-Point Store Instructions ..... 145
4.6.4 Floating-Point Load and Store Dou- ble Pair Instructions [Phased-Out] ..... 149
4.6.5 Floating-Point Move Instructions ..... 151
4.6.6 Floating-Point Arithmetic Instructions153
4.6.6.1 Floating-Point Elementary Arith- metic Instructions ..... 153
4.6.6.2 Floating-Point Multiply-Add Instruc-tions158
4.6.7 Floating-Point Rounding and Conver-sion Instructions160
4.6.7.1 Floating-Point Rounding Instruction160
4.6.7.2 Floating-Point Convert To/From Integer Instructions ..... 160
4.6.7.3 Floating Round to Integer Instruc- tions . . . . . . . . . . . . . . . . . . . . . . . . 1664.6.8 Floating-Point Compare Instructions168
4.6.9 Floating-Point Select Instruction 169
4.6.10 Floating-Point Status and ControlRegister Instructions171
Chapter 5. Decimal Floating-Point . . 175
5.1 Decimal Floating-Point (DFP) FacilityOverview .175
5.2 DFP Register Handling ..... 176
5.2.1 DFP Usage of Floating-Point Regis- ters ..... 176
5.3 DFP Support for Non-DFP Data Types178
5.4 DFP Number Representation ..... 179
5.4.1 DFP Data Format. ..... 179
5.4.1.1 Fields Within the Data Format ..... 179
5.4.1.2 Summary of DFP Data Formats 180
5.4.1.3 Preferred DPD Encoding ..... 181
5.4.2 Classes of DFP Data ..... 181
5.5 DFP Execution Model. ..... 182
5.5.1 Rounding ..... 182
5.5.2 Rounding Mode Specification ..... 183
5.5.3 Formation of Final Result. ..... 184
5.5.3.1 Use of Ideal Exponent ..... 184
5.5.4 Arithmetic Operations ..... 184
5.5.4.1 Sign of Arithmetic Result ..... 184
5.5.5 Compare Operations ..... 185
5.5.6 Test Operations ..... 185
5.5.7 Quantum Adjustment Operations 185
5.5.8 Conversion Operations ..... 185
5.5.8.1 Data-Format Conversion ..... 185
5.5.8.2 Data-Type Conversion ..... 186
5.5.9 Format Operations ..... 186
5.5.10 DFP Exceptions ..... 186
5.5.10.1 Invalid Operation Exception ..... 188
5.5.10.2 Zero Divide Exception ..... 189
5.5.10.3 Overflow Exception. ..... 189
5.5.10.4 Underflow Exception ..... 190
5.5.10.5 Inexact Exception ..... 191
5.5.11 Summary of Normal Rounding And
Range Actions. ..... 192
5.6 DFP Instruction Descriptions ..... 194
5.6.1 DFP Arithmetic Instructions ..... 195
5.6.2 DFP Compare Instructions ..... 199
5.6.3 DFP Test Instructions ..... 202
5.6.4 DFP Quantum Adjustment Instruc-
tions. ..... 205
5.6.5 DFP Conversion Instructions ..... 214
5.6.5.1 DFP Data-Format ConversionInstructions214
5.6.5.2 DFP Data-Type Conversion Instructions ..... 217
5.6.6 DFP Format Instructions ..... 219
5.6.7 DFP Instruction Summary ..... 223
Chapter 6. Vector Facility ..... 225
6.1 Vector Facility Overview ..... 225
6.2 Chapter Conventions. ..... 225
6.2.1 Description of Instruction Operation.225
6.3 Vector Facility Registers ..... 234
6.3.1 Vector Registers ..... 234
6.3.2 Vector Status and Control Register . 234
6.3.3 VR Save Register. ..... 235
6.4 Vector Storage Access Operations 2366.4.1 Accessing Unaligned Storage Oper-ands.238
6.5 Vector Integer Operations ..... 239
6.5.1 Integer Saturation ..... 239
6.6 Vector Floating-Point Operations. ..... 241
6.6.1 Floating-Point Overview ..... 241
6.6.2 Floating-Point Exceptions. ..... 241
6.6.2.1 NaN Operand Exception ..... 241
6.6.2.2 Invalid Operation Exception ..... 242
6.6.2.3 Zero Divide Exception ..... 242
6.6.2.4 Log of Zero Exception ..... 242
6.6.2.5 Overflow Exception ..... 242
6.6.2.6 Underflow Exception ..... 242
6.7 Vector Storage Access Instructions243
6.7.1 Storage Access Exceptions ..... 243
6.7.2 Vector Load Instructions. ..... 244
6.7.3 Vector Store Instructions ..... 247
6.7.4 Vector Alignment Support Instruc- tions ..... 249
6.8 Vector Permute and Formatting Instructions ..... 250
6.8.1 Vector Pack and Unpack Instructions250
6.8.2 Vector Merge Instructions ..... 257
6.8.3 Vector Splat Instructions ..... 260
6.8.4 Vector Permute Instruction ..... 262
6.8.5 Vector Select Instruction ..... 263
6.8.6 Vector Shift Instructions ..... 264
6.8.7 Vector Extract Element Instructions .269
6.8.8 Vector Insert Element Instructions . .270
6.9 Vector Integer Instructions ..... 271
6.9.1 Vector Integer Arithmetic Instructions271
6.9.1.1 Vector Integer Add Instructions 271
6.9.1.2 Vector Integer Subtract Instructions 277
6.9.1.3 Vector Integer Multiply Instructions 283
6.9.1.4 Vector Integer Multiply-Add/Sum Instructions . . . . . . . . . . . . . . . . . . . . . 287
6.9.1.5 Vector Integer Sum-Across Instructions . . . . . . . . . . . . . . . . . . . . . . . . . 292
6.9.1.6 Vector Integer Negate Instructions. 295
6.9.2 Vector Extend Sign Instructions . 296
6.9.2.1 Vector Integer Average Instructions 298
6.9.2.2 Vector Integer Absolute Difference Instructions
.300
6.9.2.3 Vector Integer Maximum and Minimum Instructions . . . . . . . . . . . . . . . . 302
6.9.3 Vector Integer Compare Instructions. 306
6.9.4 Vector Logical Instructions . . . . . 315
6.9.5 Vector Parity Byte Instructions . . 317
6.9.6 Vector Integer Rotate and Shift Instructions

318
6.10 Vector Floating-Point Instruction Set . 324
6.10.1 Vector Floating-Point Arithmetic Instructions

324
6.10.2 Vector Floating-Point Maximum and Minimum Instructions . . . . . . . . . . . . . . 326
6.10.3 Vector Floating-Point Rounding and Conversion Instructions . . . . . . . . . . . . 327
6.10.4 Vector Floating-Point Compare Instructions331
6.10.5 Vector Floating-Point Estimate Instructions ..... 334
6.11 Vector Exclusive-OR-based Instruc- tions ..... 336
6.11.1 Vector AES Instructions ..... 336
6.11.2 Vector SHA-256 and SHA-512 Sigma Instructions ..... 338
6.11.3 Vector Binary Polynomial Multiplica- tion Instructions ..... 339
6.11.4 Vector Permute and Exclusive-ORInstruction341
6.12 Vector Gather Instruction ..... 342
6.13 Vector Count Leading Zeros Instruc- tions ..... 343
6.14 Vector Count Trailing Zeros Instruc- tions . . . . . . . . . . . . . . . . . . . . . . . . . . 3446.14.1 Vector Count Leading/Trailing ZeroLSB Instructions . . . . . . . . . . . . . . . . . . 345
6.14.2 Vector Extract Element Instructions 346
6.15 Vector Population Count Instructions. 348
6.16 Vector Bit Permute Instruction . . 349
6.17 Decimal Integer Instructions . . . . 350
6.17.1 Decimal Integer Arithmetic Instructions . . . . . . . . . . . . . . . . . . . . . . . . . 350
6.17.2 Decimal Integer Format Conversion Instructions 352
6.17.3 Decimal Integer Sign Manipulation Instructions . . . . . . . . . . . . . . . . . . . . . 358
6.17.4 Decimal Integer Shift and Round
Instructions . . . . . . . . . . . . . . . . . 359
6.17.5 Decimal Integer Truncate Instructions . . . . . . . . . . . . . . . . . . . . . . . . 362
6.18 Vector Status and Control Register Instructions

364

## Chapter 7. Vector-Scalar Floating-Point Operations 365

7.1 Introduction. . . . . . . . . . . . . . . . . . 365
7.1.1 Overview of the Vector-Scalar Extension. . . . . . . . . . . . . . . . . . . . . . . . . . 365
7.1.1.1 Compatibility with Floating-Point and Decimal Floating-Point Operations 365
7.1.1.2 Compatibility with Vector Operations ................................ . . . 365
7.2 VSX Registers . . . . . . . . . . . . . . . 366
7.2.1 Vector-Scalar Registers . . . . . . . . 366
7.2.1.1 Floating-Point Registers . . . . . 366
7.2.1.2 Vector Registers . . . . . . . . . . . . 368
7.2.2 Floating-Point Status and Control Register.

369
7.3 VSX Operations . . . . . . . . . . . . . 374
7.3.1 VSX Floating-Point Arithmetic Overview. . . . . . . . . . . . . . . . . . . . . . . . . 374
7.3.2 VSX Floating-Point Data . . . . . . 375
7.3.2.1 Data Format . . . . . . . . . . . . . . 375
7.3.2.2 Value Representation . . . . . . . . 377
7.3.2.3 Sign of Result . . . . . . . . . . . . . 378
7.3.2.4 Normalization and Denormaliza-
tion . . . . . . . . . . . . . . . . . . . . . . . . . 379
7.3.2.5 Data Handling and Precision . 379
7.3.2.6 Rounding . . . . . . . . . . . . . . 383
7.3.3 VSX Floating-Point Execution Models ................................... . . . 386
7.3.3.1 VSX Execution Model for IEEE Operations.

386
7.3.3.2 VSX Execution Model for Multi-
ply-Add Type Instructions . . . . . . . . 387
7.4 VSX Floating-Point Exceptions . . . 389
7.4.1 Floating-Point Invalid Operation
Exception . . . . . . . . . . . . . . . . . . . 392
7.4.1.1 Definition. ..................... . . . 392
7.4.1.2 Action for VE=1. . . . . . . . . . . . 392
7.4.1.3 Action for VE=0. . . . . . . . . . . . 394
7.4.2 Floating-Point Zero Divide Exception 403
7.4.2.1 Definition. . . . . . . . . . . . . . . . . 403
7.4.2.2 Action for ZE=1 ..... 403
7.4.2.3 Action for ZE=0 ..... 404
7.4.3 Floating-Point Overflow Exception.406
7.4.3.1 Definition. ..... 406
7.4.3.2 Action for $\mathrm{OE}=1$ ..... 406
7.4.3.3 Action for $\mathrm{OE}=0$ ..... 409
7.4.4 Floating-Point Underflow Exception. 411
7.4.4.1 Definition ..... 411
7.4.4.2 Action for $\mathrm{UE}=1$ ..... 411
7.4.4.3 Action for UE=0 ..... 413
7.4.5 Floating-Point Inexact Exception ..... 416
7.4.5.1 Definition. ..... 416
7.4.5.2 Action for $\mathrm{XE}=1$ ..... 416
7.4.5.3 Action for XE=0 ..... 419
7.5 VSX Storage Access Operations ..... 422
7.5.1 Accessing Aligned Storage Oper- ands ..... 422
7.5.2 Accessing Unaligned Storage Oper- ands ..... 423
7.5.3 Storage Access Exceptions ..... 424
7.6 VSX Instruction Set ..... 425
7.6.1 VSX Instruction Set Summary ..... 425
7.6.1.1 VSX Storage Access Instructions.425
7.6.1.2 VSX Move Instructions ..... 426
7.6.1.3 VSX Floating-Point Arithmetic Instructions ..... 426
7.6.1.4 VSX Floating-Point Compare Instructions ..... 429
7.6.1.5 VSX FP-FP Conversion Instruc- tions ..... 430
7.6.1.6 VSX FP-Integer Conversion Instructions ..... 430
7.6.1.7 VSX Round to Floating-Point Inte- ger Instructions ..... 432
7.6.1.8 VSX Scalar Floating-Point Support Instructions . . . . . . . . . . . . . . . . . . . . 432
7.6.1.9 VSX Logical Instructions ..... 433
7.6.1.10 VSX Permute Instructions. ..... 434
7.6.2 VSX Instruction Description Conven- tions ..... 435
7.6.2.1 VSX Instruction RTL Operators 435 7.6.2.2 VSX Instruction RTL Function
Calls ..... 436
7.6.3 VSX Instruction Descriptions ..... 481
Appendix A. Suggested Floating-Point Models. ..... 779
A. 1 Floating-Point Round to Single-Preci- sion Model ..... 779
A. 2 Floating-Point Convert to Integer Model ..... 783
A. 3 Floating-Point Convert from Integer Model ..... 786
A. 4 Floating-Point Round to Integer Model 788
Appendix B. Densely Packed Decimal ..... 791
B. 1 BCD-to-DPD Translation ..... 791
B. 2 DPD-to-BCD Translation. ..... 791
B. 3 Preferred DPD encoding. ..... 792
Appendix C. Assembler Extended Mnemonics ..... 795
C. 1 Symbols ..... 795
C. 2 Branch Mnemonics ..... 796
C.2.1 BO and BI Fields ..... 796
C.2.2 Simple Branch Mnemonics ..... 796
C.2.3 Branch Mnemonics Incorporating Conditions ..... 797
C.2.4 Branch Prediction ..... 798
C. 3 Condition Register Logical Mnemonics799
C. 4 Subtract Mnemonics ..... 799
C.4.1 Subtract Immediate ..... 799
C.4.2 Subtract ..... 799
C. 5 Compare Mnemonics ..... 800
C.5.1 Doubleword Comparisons ..... 800
C.5.2 Word Comparisons ..... 800
C. 6 Trap Mnemonics ..... 801
C. 7 Integer Select Mnemonics ..... 802
C. 8 Rotate and Shift Mnemonics. ..... 803
C.8.1 Operations on Doublewords ..... 803
C.8.2 Operations on Words. ..... 804
C. 9 Move To/From Special Purpose Regis-
ter Mnemonics ..... 805
C. 10 Miscellaneous Mnemonics ..... 806
Book II:
Power ISA Virtual Environment Architecture ..... 809
Chapter 1. Storage Model ..... 811
1.1 Definitions ..... 811
1.2 Introduction ..... 812
1.3 Virtual Storage ..... 812
1.4 Single-Copy Atomicity ..... 813
1.5 Cache Model ..... 814
1.6 Storage Control Attributes ..... 814
1.6.1 Write Through Required ..... 814
1.6.2 Caching Inhibited ..... 815
1.6.3 Memory Coherence Required ..... 815
1.6.4 Guarded ..... 815
1.6.5 Strong Access Order ..... 817
1.7 Shared Storage ..... 818
1.7.1 Storage Access Ordering ..... 818
1.7.2 Storage Ordering of Copy/Paste-Initi-ated Data Transfers820
1.7.3 Storage Ordering of I/O Accesses820
1.7.4 Atomic Update ..... 820
1.7.4.1 Reservations ..... 821
1.7.4.2 Forward Progress ..... 823
1.8 Transactions ..... 823
1.8.1 Rollback-Only Transactions ..... 825
1.9 Cluster Shared Memory ..... 825
1.10 Instruction Storage ..... 826
1.10.1 Concurrent Modification and Execu- tion of Instructions ..... 828
Chapter 2. Performance Considerations and Instruction Restart ..... 829
2.1 Performance-Optimized Instruction Sequences ..... 829
2.1.1 Load and Store Operations ..... 830
2.1.2 32-Bit Constant Generation ..... 833
2.1.3 Sign and Zero Extension ..... 833
2.1.4 Load/Store Addressing Relative to Program Counter ..... 834
2.1.5 Destructive Operation Operand Preservation ..... 835
2.2 Instruction Restart ..... 836
Chapter 3. Management of Shared Resources ..... 837
3.1 Program Priority Registers ..... 837
3.2 "or" Instruction ..... 837
Chapter 4. Storage Control Instructions ..... 839
4.1 Parameters Useful to Application Pro- grams ..... 839
4.2 Data Stream Control Register (DSCR)839
4.3 Cache Management Instructions ..... 841
4.3.1 Instruction Cache Instructions ..... 842
4.3.2 Data Cache Instructions ..... 843
4.3.2.1 Obsolete Data Cache Instructions . 854
4.3.3 "or" Instruction ..... 855
4.4 Copy-Paste Facility ..... 856
4.5 Atomic Memory Operations ..... 861
4.5.1 Load Atomic ..... 861
4.5.2 Store Atomic ..... 864
4.6 Synchronization Instructions ..... 867
4.6.1 Instruction Synchronize Instruction867
4.6.2 Load and Reserve and Store Condi- tional Instructions ..... 867
4.6.2.1 64-Bit Load and Reserve and Store Conditional Instructions ..... 873
4.6.2.2 128-bit Load and Reserve StoreConditional Instructions.875
4.6.3 Memory Barrier Instructions ..... 877
4.6.4 Wait Instruction ..... 880
Chapter 5. Transactional Memory Facility ..... 881
5.1 Transactional Memory Facility Over- view ..... 881
5.1.1 Definitions ..... 882
5.2 Transactional Memory Facility States 883
5.2.1 The TDOOMED Bit ..... 885
5.3 Transaction Failure ..... 885
5.3.1 Causes of Transaction Failure ..... 885
5.3.2 Recording of Transaction Failure ..... 888
5.3.3 Handling of Transaction Failure ..... 888
5.4 Transactional Memory Facility Regis- ters ..... 889
5.4.1 Transaction Failure Handler AddressRegister (TFHAR)889
5.4.2 Transaction EXception And Status Register (TEXASR) ..... 889
5.4.3 Transaction Failure Instruction Address Register (TFIAR) ..... 891
5.5 Transactional Facility Instructions ..... 893
Chapter 6. Time Base ..... 901
6.1 Time Base Instructions ..... 902
Chapter 7. Event-Based Branch Facility ..... 905
7.1 Event-Based Branch Overview ..... 905
7.2 Event-Based Branch Registers. ..... 906
7.2.1 Branch Event Status and Control Register. ..... 906
7.2.2 Event-Based Branch Handler Regis-ter908
7.2.3 Event-Based Branch Return Register908
7.3 Event-Based Branch Instructions ..... 909
Appendix A. Assembler Extended Mnemonics ..... 911
A. 1 Data Cache Block Touch [for Store] Mnemonics911
A. 2 Data Cache Block Flush Mnemonics911
A. 3 Or Mnemonics ..... 911
A. 4 Load and Reserve Mnemonics ..... 911
A. 5 Synchronize Mnemonics ..... 912
A. 6 Wait Mnemonics ..... 912
A. 7 Transactional Memory Instruction Mnemics ..... 912
A. 8 Move To/From Time Base Mnemonics 912
A. 9 Return From Event-Based Branch Mnemonic ..... 912
Appendix B. Programming Examples for Sharing Storage ..... 913
B. 1 Atomic Update Primitives ..... 913
B. 2 Lock Acquisition and Release, and Related Techniques. ..... 915
B.2.1 Lock Acquisition and Import Barriers 915B.2.1.1 Acquire Lock and Import SharedStorage915
B.2.1.2 Obtain Pointer and Import SharedStorage915
B.2.2 Lock Release and Export Barriers916
B.2.2.1 Export Shared Storage and Release Lock ..... 916
B.2.2.2 Export Shared Storage and Release Lock using Iwsync. ..... 916
B.2.3 Safe Fetch ..... 916
B. 3 List Insertion. ..... 917
B. 4 Notes ..... 917
B. 5 Transactional Lock Elision ..... 917
B.5.1 Enter Critical Section. ..... 918
B.5.2 Handling Busy Lock. ..... 918
B.5.3 Handling TLE Abort. ..... 918
B.5.4 TLE Exit Section Critical Path ..... 918
B.5.5 Acquisition and Release of TLE Locks. ..... 918
Book III:
Power ISA Operating Environment Architecture ..... 921
Chapter 1. Introduction ..... 923
1.1 Overview ..... 923
1.2 Document Conventions ..... 923
1.2.1 Definitions and Notation ..... 923
1.2.2 Reserved Fields. ..... 924
1.3 General Systems Overview ..... 925
1.4 Exceptions ..... 925
1.5 Synchronization ..... 925
1.5.1 Context Synchronization ..... 925
1.5.2 Execution Synchronization ..... 926
Chapter 2. Logical Partitioning (LPAR) and Thread Control ..... 927
2.1 Overview ..... 927
2.2 Logical Partitioning Control Register(LPCR).927
2.3 Hypervisor Real Mode Offset Register(HRMOR).931
2.4 Logical Partition Identification Register (LPIDR) ..... 931
2.5 Processor Compatibility Register (PCR) ..... 931
2.6 Other Hypervisor Resources ..... 941
2.7 Sharing Hypervisor Resources ..... 941
2.8 Sub-Processors. ..... 942
2.9 Thread Identification Register (TIR) .....
942
2.10 Hypervisor Interrupt Little-Endian (HILE) Bit ..... 942
Chapter 3. Branch Facility ..... 943
3.1 Branch Facility Overview ..... 943
3.2 Branch Facility Registers. ..... 943
3.2.1 Machine State Register ..... 943
3.2.2 State Transitions Associated with the Transactional Memory Facility . . . . . . . 9463.2.3 Processor Stop Status and ControlRegister (PSSCR)949
3.3 Branch Facility Instructions ..... 952
3.3.1 System Linkage Instructions ..... 952
3.3.2 Power-Saving Mode. ..... 956
3.3.2.1 Power-Saving Mode Instruction957
3.3.2.2 Entering and Exiting Power-Sav-ing Mode957
3.4 Event-Based Branch Facility and Instruction ..... 959
Chapter 4. Fixed-Point Facility ..... 961
4.1 Fixed-Point Facility Overview ..... 961
4.2 Special Purpose Registers ..... 961
4.3 Fixed-Point Facility Registers ..... 961
4.3.1 Processor Version Register ..... 961
4.3.2 Chip Information Register ..... 961
4.3.3 Processor Identification Register ..... 961
4.3.4 Process Identification Register. ..... 962
4.3.5 Control Register ..... 962
4.3.6 Program Priority Register. ..... 962
4.3.7 Problem State Priority Boost Regis-
ter ..... 963
4.3.8 Relative Priority Register ..... 963
4.3.9 Software-use SPRs ..... 963
4.4 Fixed-Point Facility Instructions. ..... 965
4.4.1 Fixed-Point Load and Store Caching Inhibited Instructions ..... 965
4.4.2 OR Instruction ..... 968
4.4.3 Load Monitored Doubleword Instruc-tion . . . . . . . . . . . . . . . . . . . . . . . . . . . . 9684.4.4 Transactional Memory Instructions969
4.4.5 Move To/From System Register Instructions ..... 970
Chapter 5. Storage Control ..... 981
5.1 Overview ..... 981
5.2 Storage Exceptions ..... 981
5.3 Instruction Fetch ..... 981
5.3.1 Implicit Branch ..... 981
5.3.2 Address Wrapping Combined with Changing MSR Bit SF ..... 981
5.4 Data Access ..... 981
5.5 Performing Operations
Out-of-Order ..... 982
5.6 Invalid Real Address ..... 982
5.7 Storage Addressing ..... 983
5.7.1 32-Bit Mode. ..... 983
5.7.2 Virtualized Partition Memory (VPM) Mode. ..... 983
5.7.3 Hypervisor Real And Virtual Real Addressing Modes ..... 984
5.7.3.1 Hypervisor Offset Real Mode Address ..... 984
5.7.3.2 Storage Control Attributes for Accesses in Hypervisor Real Addressing Mode. ..... 984
5.7.3.2.1 Hypervisor Real Mode Storage Control ..... 985
5.7.3.3 Virtual Real Mode Addressing Mechanism ..... 985
5.7.3.4 Storage Control Attributes for Implicit Storage Accesses ..... 986
5.7.4 Definitions ..... 986
5.7.5 Address Ranges Having Defined Uses ..... 987
5.7.5.1 Effective Address Space Structure for Radix-using Partitions ..... 988
5.7.6 In-Memory Tables ..... 989
5.7.6.1 Partition Table ..... 989
5.7.6.2 Process Table ..... 991
5.7.7 Address Translation Overview. ..... 991
5.7.8 Segment Translation ..... 994
5.7.8.1 Segment Lookaside Buffer (SLB) . ..... 994
5.7.8.2 SLB Search ..... 996
5.7.8.3 Segment Table Description and Search ..... 996
5.7.8.3.1 Primary Hash for 256MB Seg- ment ..... 996
5.7.8.3.2 Primary Hash for 1TB Segment 996
5.7.8.3.3 Secondary Hash for 256MB Seg-ment996
5.7.8.3.4 Secondary Hash for 1TB Seg-ment996
5.7.9 Hashed Page Table Translation ..... 997
5.7.9.1 Hashed Page Table ..... 998
5.7.9.2 Page Table Search ..... 999
5.7.10 Radix Tree Translation ..... 1001
5.7.10.1 Radix Tree Page Directory Entry1002
5.7.10.2 Radix Tree Page Table Entry 1003
5.7.11 Nested Translation ..... 1003
5.7.12 Translation Process ..... 1005
5.7.12.1 Fully-Qualified Address ..... 1005
5.7.12.2 Finding the Page Tables ..... 1005
5.7.12.3 Obtaining Host Real Address,Radix on Radix1006
5.7.12.4 Obtaining Host Real Address, HPT ..... 1007
5.7.13 Reference and Change Recording10075.7.14 Storage Protection1011
5.7.14.1 Virtual Page Class Key Protection1011
5.7.14.2 Basic Storage Protection, Address Translation Enabled ..... 1015
5.7.14.3 Basic Storage Protection, Address Translation Disabled ..... 1016
5.7.14.4 Radix Tree Translation Storage Protection ..... 1016
5.7.15 Cluster Shared Memory Protection1017
5.8 Storage Control Attributes ..... 1017
5.8.1 Guarded Storage ..... 1017
5.8.1.1 Out-of-Order Accesses to Guarded Storage ..... 1018
5.8.2 Storage Control Bits. ..... 1018
5.8.2.1 Storage Control Bit Restrictions1019
5.8.2.2 Altering the Storage Control Bits 1019
5.9 Storage Control Instructions. ..... 1021
5.9.1 Cache Management Instructions1021
5.9.2 Synchronize Instruction ..... 1021
5.9.3 Lookaside Buffer Management ..... 1022
5.9.3.1 Thread-Specific Segment Transla- tions ..... 1023
5.9.3.2 SLB Management Instructions ..... 1023
5.9.3.3 TLB Management Instructions1033
5.10 Translation Table Update Synchroni-zation Requirements1043
5.10.1 Translation Table Updates ..... 1043
5.10.1.1 Adding a Page Table Entry . ..... 1045
5.10.1.2 Modifying a Translation Table Entry ..... 1045
Chapter 6. Interrupts ..... 1047
6.1 Overview ..... 1047
6.2 Interrupt Registers ..... 1047
6.2.1 Machine Status Save/Restore Regis- ters . . . . . . . . . . . . . . . . . . . . . . . . . . 1047
6.2.2 Hypervisor Machine Status Save/
Restore Registers ..... 1047
6.2.3 Access Segment Descriptor Register 1047
6.2.4 Data Address Register ..... 1048
6.2.5 Hypervisor Data Address Register1048
6.2.6 Data Storage Interrupt Status Register ..... 1048
6.2.7 Hypervisor Data Storage Interrupt Status Register ..... 1048
6.2.8 Hypervisor Emulation InstructionRegister.1048
6.2.9 Hypervisor Maintenance Exception Register. ..... 1049
6.2.10 Hypervisor Maintenance Exception Enable Register. ..... 1049
6.2.11 Facility Status and Control Register1049
6.2.12 Hypervisor Facility Status and Con-trol Register. . . . . . . . . . . . . . . . . . . . 1051
6.3 Interrupt Synchronization ..... 1055
6.4 Interrupt Classes ..... 1055
6.4.1 Precise Interrupt ..... 1055
6.4.2 Imprecise Interrupt ..... 1055
6.4.3 Interrupt Processing ..... 1057
6.4.4 Implicit alteration of HSRRO and HSRR1 ..... 1059
6.5 Interrupt Definitions ..... 1060
6.5.1 System Reset Interrupt ..... 1062
6.5.2 Machine Check Interrupt ..... 1064
6.5.3 Data Storage Interrupt ..... 1065
6.5.4 Data Segment Interrupt ..... 1067
6.5.5 Instruction Storage Interrupt ..... 1068
6.5.6 Instruction Segment Interrupt. ..... 1069
6.5.7 External Interrupt ..... 1069
6.5.7.1 Direct External Interrupt ..... 1069
6.5.7.2 Mediated External Interrupt ..... 1070
6.5.8 Alignment Interrupt ..... 1070
6.5.9 Program Interrupt. ..... 1071
6.5.10 Floating-Point Unavailable Interrupt. ..... 1073
6.5.11 Decrementer Interrupt ..... 1073
6.5.12 Hypervisor Decrementer Interrupt. ..... 1073
6.5.13 Directed Privileged Doorbell Inter- rupt ..... 1073
6.5.14 System Call Interrupt ..... 1074
6.5.15 Trace Interrupt ..... 1074
6.5.16 Hypervisor Data Storage Interrupt .1075
6.5.17 Hypervisor Instruction StorageInterrupt1078
6.5.18 Hypervisor Emulation Assistance Interrupt ..... 1080
6.5.19 Hypervisor Maintenance Interrupt .1081
6.5.20 Directed Hypervisor Doorbell Inter-rupt1081
6.5.21 Hypervisor Virtualization Interrupt .1081
6.5.22 Performance Monitor
Interrupt ..... 1082
6.5.23 Vector Unavailable Interrupt . ..... 1082
6.5.24 VSX Unavailable Interrupt ..... 1082
6.5.25 Facility Unavailable Interrupt ..... 1082
6.5.26 Hypervisor Facility Unavailable Interrupt ..... 1083
6.5.27 System Call Vectored Interrupt1083
6.6 Partially ExecutedInstructions1084
6.7 Exception Ordering ..... 1085
6.7.1 Unordered Exceptions ..... 1085
6.7.2 Ordered Exceptions ..... 1085
6.8 Event-Based Branch Exception Order- ing ..... 1086
6.9 Interrupt Priorities ..... 1086
6.10 Relationship of Event-Based
Branches to Interrupts ..... 1089
6.10.1 EBB Exception Priority ..... 1089
6.10.2 EBB Synchronization ..... 1089
6.10.3 EBB Classes ..... 1089
Chapter 7. Timer Facilities ..... 1091
7.1 Overview ..... 1091
7.2 Time Base (TB) ..... 1091
7.2.1 Writing the Time Base ..... 1092
7.3 Virtual Time Base ..... 1092
7.4 Decrementer ..... 1093
7.4.1 Writing and Reading the Decre- menter ..... 1094
7.5 Hypervisor Decrementer ..... 1094
7.6 Processor Utilization of ResourcesRegister (PURR)1094
7.7 Scaled Processor Utilization of Resources Register (SPURR) ..... 1095
7.8 Instruction Counter ..... 1096
Chapter 8. Debug Facilities ..... 1097
8.1 Overview ..... 1097
8.2 Come-From Address Register. ..... 1097
8.3 Completed Instruction Address Break-point.1097
8.4 Data Address Watchpoint ..... 1098
Chapter 9. Performance Monitor Facility ..... 1101
9.1 Overview ..... 1101
9.2 Performance Monitor Operation. ..... 1101
9.3 No-op Instructions Reserved for the
Performance Monitor ..... 1102
9.4 Performance Monitor Facility Registers 1102
9.4.1 Performance Monitor SPR Numbers.
1102
9.4.2 Performance Monitor Counters. 1103
9.4.2.1 Event Counting and Sampling ..... 1103
9.4.3 Threshold Event Counter ..... 1104
9.4.4 Monitor Mode Control Register 0 ..... 1105
9.4.5 Monitor Mode Control Register 11110
9.4.6 Monitor Mode Control Register 2 ..... 1113
9.4.7 Monitor Mode Control Register A ..... 1113
9.4.8 Sampled Instruction Address Regis-
ter. ..... 1116
9.4.9 Sampled Data Address Register1117
9.4.10 Sampled Instruction Event Register1117
9.5 Branch History Rolling Buffer ..... 1119
9.5.1 BHRB Filtering ..... 1119
9.6 Interaction With Other Facilities ..... 1120
Chapter 10. Processor Control . 1121
10.1 Overview ..... 1121
10.2 Programming Model ..... 1121
10.3 Processor Control Registers ..... 1121
10.3.1 Directed Privileged Doorbell Excep-tion State1121
10.4 Processor Control Instructions. ..... 1123
Chapter 11. SynchronizationRequirements for Context Alterations1127
Power ISA Book I-III Appendices ..... 1133
Appendix A. Illegal Instructions ..... 1135
Appendix B. Reserved Instructions.1137
Appendix C. Opcode Maps ..... 1139
Appendix D. Power ISA InstructionSet Sorted by Opcode1173
Appendix E. Power ISA Instruction Set Sorted by Version ..... 1191
Appendix F. Power ISA Instruction Set Sorted by Mnemonic ..... 1209
Last Page - End of Document ..... 1227

## Book I:

## Power ISA User Instruction Set Architecture

Version 3.0

## Chapter 1. Introduction

### 1.1 Overview

This chapter describes computation modes, document conventions, a processor overview, instruction formats, storage addressing, and instruction fetching.

### 1.2 Instruction Mnemonics and Operands

The description of each instruction includes the mnemonic and a formatted list of operands. Some examples are the following.

```
stw RS,D(RA)
addis RT,RA,SI
```

Power ISA-compliant Assemblers will support the mnemonics and operand lists exactly as shown. They should also provide certain extended mnemonics, such as the ones described in Appendix C of Book I.

### 1.3 Document Conventions

### 1.3.1 Definitions

The following definitions are used throughout this document.

- program

A sequence of related instructions.

- application program

A program that uses only the instructions and resources described in Books I and II.

- processor

The hardware component that implements the instruction set, storage model, and other facilities defined in the Power ISA architecture, and executes the instructions specified in a program.

- quadword, doubleword, word, halfword, and byte
128 bits, 64 bits, 32 bits, 16 bits, and 8 bits, respectively.
- positive

Means greater than zero.

- negative

Means less than zero.
■ floating-point single format (or simply single format)
Refers to the representation of a single-precision binary floating-point value in a register or storage.

■ floating-point double format (or simply double format)
Refers to the representation of a double-precision
binary floating-point value in a register or storage.
■ system library program
A component of the system software that can be called by an application program using a Branch instruction.

- system service program

A component of the system software that can be called by an application program using a System Call or System Call Vectored instruction.

- system trap handler

A component of the system software that receives control when the conditions specified in a Trap instruction are satisfied.
■ system error handler
A component of the system software that receives control when an error occurs. The system error handler includes a component for each of the various kinds of error. These error-specific components are referred to as the system alignment error handler, the system data storage error handler, etc.

■ latency
Refers to the interval from the time an instruction begins execution until it produces a result that is available for use by a subsequent instruction.

■ unavailable
Refers to a resource that cannot be used by the program. For example, storage is unavailable if access to it is denied. See Book III.

## ■ undefined value

May vary between implementations, and between different executions on the same implementation, and similarly for register contents, storage contents, etc., that are specified as being undefined.

## - boundedly undefined

The results of executing a given instruction are said to be boundedly undefined if they could have been achieved by executing an arbitrary finite sequence of instructions (none of which yields boundedly undefined results) in the state the processor was in before executing the given instruction. Boundedly undefined results may include the presentation of inconsistent state to the system error handler as described in Section 1.10.1 of Book II. Boundedly undefined results for a given instruction may vary between implementations, and between different executions on the same implementation.

## ■ "must"

If software violates a rule that is stated using the word "must" (e.g., "this field must be set to 0"), the results are boundedly undefined unless otherwise stated.

## - sequential execution model

The model of program execution described in Section 2.2, "Instruction Execution Order" on page 29.

### 1.3.2 Notation

The following notation is used throughout the Power ISA documents.

- All numbers are decimal unless specified in some special way.
- Obnnnn means a number expressed in binary format.
- Oxnnnn means a number expressed in hexadecimal format.
Underscores may be used between digits.
- RT, RA, R1, ... refer to General Purpose Registers.
- FRT, FRA, FR1, ... refer to Floating-Point Registers.
- FRTp, FRAp, FRBp, ... refer to an even-odd pair of Floating-Point Registers. Values must be even, otherwise the instruction form is invalid.
- VRT, VRA, VR1, ... refer to Vector Registers.
- ( x ) means the contents of register x , where x is the name of an instruction field. For example, (RA) means the contents of register RA, and (FRA) means the contents of register FRA, where RA and FRA are instruction fields. Names such as LR and CTR denote registers, not fields, so parenthe-
ses are not used with them. Parentheses are also omitted when register x is the register into which the result of an operation is placed.
- (RAIO) means the contents of register RA if the RA field has the value 1-31, or the value 0 if the RA field is 0 .
■ Bytes in instructions, fields, and bit strings are numbered from left to right, starting with byte 0 (most significant).
■ Bits in registers, instructions, fields, and bit strings are specified as follows. In the last three items (definition of $X_{p}$ etc.), if $X$ is a field that specifies a GPR, FPR, or VR (e.g., the RS field of an instruction), the definitions apply to the register, not to the field.
- Bits in instructions, fields, and bit strings are numbered from left to right, starting with bit 0
- For all registers except the Vector registers, bits in registers that are less than 64 bits start with bit number 64-L, where $L$ is the register length; for the Vector registers, bits in registers that are less than 128 bits start with bit number 128-L.
- The leftmost bit of a sequence of bits is the most significant bit of the sequence.
- $\quad X_{p}$ means bit $p$ of register/instruction/field/ bit_string $X$.
- $\quad X_{p: q}$ means bits $p$ through $q$ of register/instruction/field/bit_string $X$.
- $\quad X_{p q \ldots}$ means bits $p, q, \ldots$ of register/instruction/field/bit_string $X$.
- $\neg(R A)$ means the one's complement of the contents of register RA.
- A period (.) as the last character of an instruction mnemonic means that the instruction records status information in certain fields of the Condition Register as a side effect of execution.
- The symbol II is used to describe the concatenation of two values. For example, 010 \|I 111 is the same as 010111.
- $x^{n}$ means $x$ raised to the $n^{\text {th }}$ power.
- ${ }^{n} x$ means the replication of $x, n$ times (i.e., $x$ concatenated to itself $n-1$ times). ${ }^{n} 0$ and ${ }^{n} 1$ are special cases:
- ${ }^{n} 0$ means a field of $n$ bits with each bit equal to 0 . Thus ${ }^{5} 0$ is equivalent to $0 b 00000$.
- $\quad{ }^{n} 1$ means a field of $n$ bits with each bit equal to 1. Thus ${ }^{5} 1$ is equivalent to $0 b 11111$.
- Each bit and field in instructions, and in status and control registers (e.g., XER, FPSCR) and Special Purpose Registers, is either defined or reserved. Some defined fields contain reserved values. In such cases when this document refers to the specific field, it refers only to the defined values, unless otherwise specified.

■ /, I/, I/I, ... denotes a reserved field, in a register, instruction, field, or bit string.
■ ?, ??, ???, ... denotes an implementation-dependent field in a register, instruction, field or bit string.

### 1.3.3 Reserved Fields, Reserved Values, and Reserved SPRs

Reserved fields in instructions are ignored by the processor.

In some cases a defined field of an instruction has certain values that are reserved. This includes cases in which the field is shown in the instruction layout as containing a particular value; in such cases all other values of the field are reserved. In general, if an instruction is coded such that a defined field contains a reserved value the instruction form is invalid; see Section 1.9.2 on page 23. The only exception to the preceding rule is that it does not apply to Reserved and IIlegal classes of instructions (see Section 1.8) or to portions of defined fields that are specified, in the instruction description, as being treated as reserved fields.
To maximize compatibility with future architecture extensions, software must ensure that reserved fields in instructions contain zero and that defined fields of instructions do not contain reserved values.

The handling of reserved bits in System Registers (e.g., XER, FPSCR) depends on whether the processor is in problem state. Unless otherwise stated, software is permitted to write any value to such a bit. In problem state, a subsequent reading of the bit returns 0 regardless of the value written; in privileged states, a subsequent reading of the bit returns 0 if the value last written to the bit was 0 and returns an undefined value ( 0 or 1 ) otherwise.

In some cases, a defined field of a System Register has certain values that are reserved. Software must not set a defined field of a System Register to a reserved value. References elsewhere in this document to a defined field (in an instruction or System Register) that has reserved values assume the field does not contain a reserved value, unless otherwise stated or obvious from context.

In some cases, a given bit of a System Register is specified to be set to a constant value by a given instruction or event. Unless otherwise stated or obvious from context, software should not depend on this constant value because the bit may be assigned a meaning in a future version of the architecture.

The reserved SPRs include SPRs 808, 809, 810, and 811. mtspr and mfspr instructions specifying these SPRs are treated as noops. Reserved SPRs are provided in the architecture to anticipate the eventual adoption of performance hint functionality that must be controlled by SPRs. Control of these capabilities using reserved SPRs will allow software to use these new capabilities on new implementations that support them while remaining compatible with existing implementations that may not support the new functionality.

Reserved SPRs are not assigned names. There are no individual descriptions of reserved SPRs in this document.

## Assembler Note

Assemblers should report uses of reserved values of defined fields of instructions as errors.

## Programming Note

It is the responsibility of software to preserve bits that are now reserved in System Registers, because they may be assigned a meaning in some future version of the architecture.
In order to accomplish this preservation in imple-mentation-independent fashion, software should do the following.

- Initialize each such register supplying zeros for all reserved bits.
- Alter (defined) bit(s) in the register by reading the register, altering only the desired bit(s), and then writing the new value back to the register.
The XER and FPSCR are partial exceptions to this recommendation. Software can alter the status bits in these registers, preserving the reserved bits, by executing instructions that have the side effect of altering the status bits. Similarly, software can alter any defined bit in the FPSCR by executing a Float-ing-Point Status and Control Register instruction. Using such instructions is likely to yield better performance than using the method described in the second item above.


### 1.3.4 Description of Instruction Operation

Instruction descriptions (including related material such as the introduction to the section describing the instructions) mention that the instruction may cause a system error handler to be invoked, under certain conditions, if and only if the system error handler may treat the case as a programming error. (An instruction may cause a system error handler to be invoked under other conditions as well; see Chapter 6 of Book III).
A formal description is given of the operation of each instruction. In addition, the operation of most instructions is described by a semiformal language at the register transfer level (RTL). This RTL uses the notation given below, in addition to the notation described in Section 1.3.2. Some of this notation is also used in the formal descriptions of instructions. RTL notation not summarized here should be self-explanatory.
The RTL descriptions cover the normal execution of the instruction, except that "standard" setting of status registers, such as the Condition Register, is not shown.
("Non-standard" setting of these registers, such as the setting of the Condition Register by the Compare instructions, is shown.) The RTL descriptions do not cover cases in which the system error handler is invoked, or for which the results are boundedly undefined.
The RTL descriptions specify the architectural transformation performed by the execution of an instruction. They do not imply any particular implementation.

## Notation Meaning

| $\leftarrow$ | Assignment |
| :---: | :---: |
| $\leftarrow_{i e a}$ | Assignment of an instruction effective address. In 32-bit mode the high-order 32 bits of the 64-bit target address are set to 0. |
| ᄀ | NOT logical operator |
| + | Two's complement addition |
| - | Two's complement subtraction, unary minus |
| $\times$ | Multiplication |
| $\times_{\text {si }}$ | Signed-integer multiplication |
| $\times_{\text {ui }}$ | Unsigned-integer multiplication |
| / | Division |
| $\div$ | Division, with result truncated to integer |
| \% | Remainder of integer division |
| $\checkmark$ | Square root |
| =, $\neq$ | Equals, Not Equals relations |
| <, $\leq,>, \geq$ | Signed comparison relations |
| $<^{\text {u }}$, > ${ }^{\text {u }}$ | Unsigned comparison relations |
| ? | Unordered comparison relation |
| \& I | AND, OR logical operators |
| $\oplus$, | Exclusive OR, Equivalence logical operators $((a \equiv b)=(a \oplus \neg b))$ |

$\operatorname{ABS}(x) \quad$ Absolute value of $x$
BCD_TO_DPD(x)
The low-order 24 bits of $x$ contain six, 4-bit BCD fields which are converted to two declets; each set of two declets is placed into the low-order 20 bits of the result. See Section B.1, "BCD-to-DPD Translation".
CEIL(x) Least integer $\geq x$
DOUBLE(x) Result of converting $x$ from floating-point single format to floating-point double format, using the model shown on page 141
DPD_TO_BCD(x)
The low-order 20 bits of $x$ contain two declets which are converted to six, 4-bit BCD fields; each set of six, 4-bit BCD fields is placed into the low-order 24 bits of the result. See Section B.2, "DPD-to-BCD Translation".
EXTS( x ) Result of extending $x$ on the left with sign bits
FLOOR(x) Greatest integer $\leq x$
GPR(x) General Purpose Register $x$
MASK( $\mathrm{x}, \mathrm{y}$ ) Mask having 1 s in positions x through y (wrapping if $x>y$ ) and Os elsewhere

MEM( $x, y$ ) Contents of a sequence of $y$ bytes of storage. The sequence depends on the byte ordering used for storage access, as follows.
Big-Endian byte ordering:
The sequence starts with the byte at address $x$ and ends with the byte at address $x+y-1$.
Little-Endian byte ordering:
The sequence starts with the byte at address $x+y-1$ and ends with the byte at address $x$.
$\operatorname{ROTL}_{64}(\mathrm{x}, \mathrm{y})$
Result of rotating the 64-bit value x left y positions
ROTL $_{32}(x, y)$
Result of rotating the 64-bit value xllx left y positions, where $x$ is 32 bits long
SINGLE(x) Result of converting $x$ from floating-point double format to floating-point single format, using the model shown on page 145
SPR(x) Special Purpose Register $x$
switch/case/default
switch/case/default statement, indenting shows range. The clause after "switch" specifies the expression to evaluate. The clause after "case" specifies individual values for the expression, followed by a colon, followed by the actions that are taken if the evaluated expression has any of the specified values. "default" is optional. If present, it must follow all the "case" clauses. The clause after "default" starts with a colon, and specifies the actions that are taken if the evaluated expression does not have any of the values specified in the preceding case statements.
TRAP Invoke the system trap handler
characterization
Reference to the setting of status bits, in a standard way that is explained in the text
undefined An undefined value.
CIA Current Instruction Address, which is the 64-bit address of the instruction being described by a sequence of RTL. Used by relative branches to set the Next Instruction Address (NIA), and by Branch instructions with LK=1 to set the Link Register. Does not correspond to any architected register. The CIA is sometimes referred to as the Program Counter (PC).
NIA Next Instruction Address, which is the 64-bit address of the next instruction to be executed. For a successful branch, the next instruction address is the branch target address: in RTL, this is indicated by assigning a value to NIA. For other instructions that cause non-sequential instruction fetching (see Book III), the RTL is similar.

For instructions that do not branch, and do not otherwise cause instruction fetching to be non-sequential, the next instruction address is CIA+4. Does not correspond to any architected register.
if... then... else...
Conditional execution, indenting shows range; else is optional.
do
Do loop, indenting shows range. "To" and/ or "by" clauses specify incrementing an iteration variable, and a "while" clause gives termination conditions.
leave Leave innermost do loop, or do loop described in leave statement.
for $\quad$ For loop, indenting shows range. Clause after "for" specifies the entities for which to execute the body of the loop.

## Version 3.0

The precedence rules for RTL operators are summarized in Table 1. Operators higher in the table are applied before those lower in the table. Operators at the same level in the table associate from left to right, from right to left, or not at all, as shown. (For example, associates from left to right, so $a-b-c=(a-b)-c$.) Parentheses are used to override the evaluation order implied by the table or to increase clarity; parenthesized expressions are evaluated before serving as operands.

| Table 1: Operator precedence |  |
| :--- | :---: |
| Operators | Associativity |
| subscript, function evaluation | left to right |
| pre-superscript (replication), <br> post-superscript (exponentiation) | right to left |
| unary,-$\urcorner$ | right to left |
| $\times, \div$ | left to right |
| ,,+- | left to right |
| II | left to right |
| $=, \neq,<, \leq,>, \geq,<^{u},>^{\text {u }}, ?$ | left to right |
| $\&, \oplus, \equiv$ | left to right |
| l | left to right |
| $:($ range $)$ | none |
| $\leftarrow, \leftarrow$ iea | none |

### 1.3.5 Phased-Out Facilities

## Phased-Out Facilities

These are facilities and instructions that, in some future version of the architecture, will be dropped out of the architecture. System developers should develop a migration plan to eliminate use of them in new systems. These facilities are marked with a [Phased-Out] marker.

Phased-Out facilities and instructions must be implemented.

## Programming Note

Warning: Instructions and facilities being phased out of the architecture are likely to perform poorly on future implementations. New programs should not use them.

### 1.4 Processor Overview

The basic classes of instructions are as follows:

- branch instructions (Chapter 2)
- GPR-based scalar fixed-point instructions (Chapter 3)
■ FPR-based scalar floating-point instructions (Chapter 4)
- FPR-based scalar decimal floating-point instructions (Chapter 5)
■ VR-based vector fixed-point and floating-point instructions (Chapter 6)
■ VSR-based scalar and vector floating-point instructions (Chapter 7)

Scalar fixed-point instructions operate on byte, halfword, word, doubleword, and quadword operands, where each operand contained in a GPR. Vector fixed-point instructions operate on vectors of byte, halfword, and word operands, where each vector is con-
| tained in a VR. Scalar floating-point instructions operate on single-precision or double-precision float-ing-point operands, where each operand is contained
I in an FPR or VSR. Vector floating-point instructions operate on vectors of single-precision and double-precision floating-point operands, where each vector is
I contained in a VR or VSR.
The Power ISA uses instructions that are four bytes long and word-aligned. It provides for byte, halfword, word, doubleword, and quadword operand loads and stores between storage and a set of 32 General Purpose Registers (GPRs). It provides for word and doubleword operand loads and stores between storage and a set of 32 Floating-Point Registers (FPRs). It also provides for byte, halfword, word, and quadword operand loads and stores between storage and a set of 32 Vector Registers (VRs). It provides for doubleword and quadword operand loads and stores between storage and a set of 64 Vector-Scalar Registers (VSRs).

Signed integers are represented in two's complement form.

There are no computational instructions that modify storage; instructions that reference storage may reformat the data (e.g. load halfword algebraic). To use a storage operand in a computation and then modify the same or another storage location, the contents of the storage operand must be loaded into a register, modified, and then stored back to the target location. Figure 1 is a logical representation of instruction processing. Figure 2 shows the registers that are defined in Book I. (A few additional registers that are available to application programs are defined in other Books, and are not shown in the figure.)


Figure 1. Logical processing model

## Version 3.0


"Condition Register" on page 30

| LR |  |
| :--- | :--- |
| 0 | 63 |
| "Link Register" on page 32 |  |


|  | CTR |
| :--- | :--- |
| 0 | 63 |

"Count Register" on page 32

| GPR 0 |
| :---: |
| GPR 1 |
| $\cdots$ |
| $\cdots$ |
| GPR 30 |
| GPR 31 |

"General Purpose Registers" on page 45

|  |
| :--- |
| 0 |

"Fixed-Point Exception Register" on page 45

"VR Save Register" on page 235

| FPR 0 |  |
| :---: | :---: |
| FPR 1 |  |
| $\cdots$ |  |
| $\cdots$ |  |
| FPR 30 |  |
| FPR 31 |  |
| 0 |  |

"Floating-Point Registers" on page 124

| FPSCR |  |  |  |
| :--- | :--- | :--- | :---: |
| 32 <br> "Floating-Point <br> page 124 | Status and Control Register" on |  |  |


| VR 0 |
| :---: |
| VR 1 |
| $\ldots$ |
| $\ldots$ |
| VR 30 |
| VR 31 |
| 0 |

"Vector Registers" on page 234

| VSCR |  |  |
| :--- | :--- | :---: |
| 96 |  |  |

"Vector Status and Control Register" on page 234

| VSR 0 |
| :---: |
| VSR 1 |
| $\ldots$ |
| $\ldots$ |
| VSR 62 |
| 0 |

"Vector-Scalar Registers" on page 366

Figure 2. Registers that are defined in Book I

### 1.5 Computation modes

Processors provide two execution modes, 64-bit mode and 32 -bit mode. In both of these modes, instructions that set a 64-bit register affect all 64 bits. The computational mode controls how the effective address is interpreted, how Condition Register bits and XER bits are set, how the Link Register is set by Branch instructions
in which LK=1, and how the Count Register is tested by Branch Conditional instructions. Nearly all instructions are available in both modes (the only exceptions are a few instructions that are defined in Book III). In both modes, effective address computations use all 64 bits of the relevant registers (General Purpose Registers,

Link Register, Count Register, etc.) and produce a 64 -bit result. However, in 32-bit mode the high-order 32 bits of the computed effective address are ignored for the purpose of addressing storage; see Section 1.11.3 for additional details.

## Programming Note

Although instructions that set a 64-bit register affect all 64 bits in both 32-bit and 64-bit modes, operating systems often do not preserve the upper 32-bits of all registers across context switches done in 32-bit mode. For this reason, application programs operating in 32-bit mode should not assume that the upper 32 bits of the GPRs are preserved from instruction to instruction unless the operating system is known to preserve these bits.

### 1.6 Instruction Formats

All instructions are four bytes long and word-aligned. Thus, whenever instruction addresses are presented to the processor (as in Branch instructions) the low-order two bits are ignored. Similarly, whenever the processor develops an instruction address the low-order two bits are zero.

Bits 0:5 always specify the primary opcode (PO, below). Many instructions also have an extended opcode (XO, below). The remaining bits of the instruction contain one or more fields as shown below for the different instruction formats.

The format diagrams given below show horizontally all valid combinations of instruction fields. The diagrams include instruction fields that are used only by instructions defined in Book II or in Book III.

## Split Field Notation

In some cases an instruction field occupies more than one contiguous sequence of bits, or occupies one contiguous sequence of bits that are used in permuted order. Such a field is called a split field. In the format diagrams given below and in the individual instruction layouts, the name of a split field is shown in small letters, once for each of the contiguous sequences. In the RTL description of an instruction having a split field, and in certain other places where individual bits of a split field are identified, the name of the field in small letters represents the concatenation of the sequences from left to right. In all other places, the name of the field is capitalized and represents the concatenation of the sequences in some order, which need not be left to right, as described for each affected instruction.

### 1.6.1 A-FORM

| 0 | 6 | 11 | 16 | 21 | 26 | 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | FRT | III | FRB | I/I | XO | RC |
| PO | FRT | FRA | III | FRC | XO | Rc |
| PO | FRT | FRA | FRB | I/I | XO | Rc |
| PO | FRT | FRA | FRB | FRC | XO | Rc |
| PO | RT | RA | RB | BC | XO | / |

Figure 3. A instruction format

### 1.6.2 B-FORM

| 6 |  | 11 | 16 | 3031 |
| :---: | :---: | :---: | :---: | :---: |
| PO | BO | BI | BD | LK |

Figure 4. $B$ instruction format

### 1.6.3 D-FORM

| 6 |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| PO | BF | $/ \mathrm{L}$ | RA | SI |
| PO | BF | $/ \mathrm{L}$ | RA | UI |
| PO | FRS | RA | D |  |
| PO | FRT | RA | D |  |
| PO | RS | RA | D |  |
| PO | RS | RA | UI |  |
| PO | RT | RA | D |  |
| PO | RT | RA | SI |  |
| PO | TO | RA | SI |  |

Figure 5. D instruction format

### 1.6.4 DQ-FORM

|  |  | 11 |  | 16 | 282931 |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | PO | RTp | RA | DQ | PT |
|  | PO | S | RA | DQ | SX XO |
| I | PO | T | RA | DQ | TX XO |

Figure 6. DQ instruction format

### 1.6.5 DS-FORM

| PO | FRSp | RA | DS | XO |
| :---: | :---: | :---: | :---: | :---: |
| PO | FRTp | RA | DS | XO |
| PO | RS | RA | DS | XO |
| PO | RSp | RA | DS | XO |
| PO | RT | RA | DS | XO |
| PO | VRS | RA | DS | XO |
| PO | VRT | RA | DS | XO |

Figure 7. DS instruction format

### 1.6.6 DX-FORM


| Figure 8. DX instruction format

### 1.6.7 I-FORM

| 6 | 3031 |  |
| :---: | :---: | :---: | :---: |
| PO | LI | AA LK |

Figure 9. I instruction format

### 1.6.8 M-FORM

| 6 | 11 | 16 | 21 | 26 |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | RS | RA | RB | MB | ME | Rc |
| PO | RS | RA | SH | MB | ME | Rc |

Figure 10. $M$ instruction format

### 1.6.9 MD-FORM

|  | 6 | 11 | 16 | 21 | 27 | 3031 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | RS | RA | sh | mb | XO | sh Rc |
| PO | RS | RA | sh | me | XO | sh RC |

Figure 11. MD instruction format

### 1.6.10 MDS-FORM



Figure 12. MDS instruction format

### 1.6.11 SC-FORM



Figure 13. SC instruction format

### 1.6.12 VA-FORM

| 6 |  | 11 | 2122 |  | 26 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| PO | RT | RA | RB | RC | XO |
| PO | VRT | VRA | VRB | $\prime$ | SHB |
| PO | VRT | VRA | VRB | VRC | XO |

Figure 14. VA instruction format

### 1.6.13 VC-FORM

| 6 |  | 11 | 16 | 2122 |  | 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | VRT | VRA | VRB | RC | XO |  |

Figure 15. VC instruction format

### 1.6.14 VX-FORM

| 0 | 6 | $11121314 \quad 16$ |  | 212223 |  | 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | III | III | VRB |  | XO |  |
| PO | RT | XO | VRB |  | XO |  |
| PO | VRT | III | III |  | XO |  |
| PO | VRT | III | VRB |  | XO |  |
| PO | VRT | III ${ }^{\text {UIIM }}$ | VRB |  | XO |  |
| PO | VRT | // ${ }^{\text {U UIM }}$ | VRB |  | XO |  |
| PO | VRT | / UIM | VRB |  | XO |  |
| PO | VRT | RA | VRB |  | XO |  |
| PO | VRT | SIM | I/I |  | XO |  |
| PO | VRT | UIM | VRB |  | XO |  |
| PO | VRT | VRA | III |  | XO |  |
| PO | VRT | VRA | VRB | $1 / 1$ | XO |  |
| PO | VRT | VRA | VRB | 1 PS | XO |  |
| PO | VRT | VRA | VRB |  | XO |  |
| PO | VRT | EO | VRB | $1 / 1$ | XO |  |
| PO | VRT | EO | VRB | 1 PS | XO |  |
| PO | VRT | EO | VRB |  | XO |  |

Figure 16. VX instruction format

### 1.6.15 X-FORM

I

| PO | III |  | III | III | XO | / |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | III |  | I/I | I/I | XO | 1 |
| PO | III |  | I/I | RB | XO | / |
| PO | III |  | RA | III | XO | / |
| PO | III |  | RA | I/I | XO | 1 |
| PO | III |  | RA | RB | XO | / |
| PO | III | L | I/I | I/I | XO | / |
| PO | III | L | I/I | RB | XO | / |
| PO | III |  | RA | RB | XO | / |
| PO | III |  | RA | RB | XO | Rc |
| PO | III | L | I/I | I/I | XO | 1 |
| PO | III | L | RA | RB | XO | , |
| PO | III | WC | I/I | I/I | XO | / |
| PO | // | IH | I/I | III | XO | / |
| PO | 1 CT |  | RA | RB | XO | / |
| PO | A III |  | III | III | XO | / |
| PO | A III R |  | III | III | XO | / |
| PO | BF | // | III | III | XO | / |
| PO | BF | // | I/I | FRB | XO | / |
| PO | BF | // | III W | U | XO | Rc |
| PO | BF | // | BFA // | I/I | XO | / |
| PO | BF | // | FRA | FRB | XO | / |

Figure 17. X instruction format

| PO | BF | // | FRA | FRBp | XO | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | BF | // | FRAp | FRBp | XO | / |
| PO | BF | // | RA | RB | XO | / |
| PO | BF | // | UIM | FRB | XO | 1 |
| PO | BF | // | UIM | FRBp | XO | 1 |
| PO | BF | // | VRA | VRB | XO | / |
| PO | BF / | 11 | RA | RB | XO | 1 |
| PO | BF / | 1 L | RA | RB | XO | 1 |
| PO | BF |  | DCMX | VRB | XO | 1 |
| PO | BT |  | III | I/I | XO | Rc |
| PO | FRS |  | RA | RB | XO | / |
| PO | FRSp |  | RA | RB | XO | 1 |
| PO | FRT |  | I/I | /I/ | XO | Rc |
| PO | FRT |  | III | FRB | XO | Rc |
| PO | FRT |  | I/I | FRBp | XO | Rc |
| PO | FRT |  | FRA | FRB | XO | / |
| PO | FRT |  | FRA | FRB | XO | Rc |
| PO | FRT |  | RA | RB | XO | / |
| PO | FRT |  | S III | FRB | XO | RC |
| PO | FRT |  | SP III | FRB | XO | Rc |
| PO | FRTp |  | III | FRB | XO | RC |
| PO | FRTp |  | I/I | FRBp | XO | Rc |
| PO | FRTp |  | FRA | FRBp | XO | Rc |
| PO | FRTp |  | FRAp | FRBp | XO | Rc |
| PO | FRTp |  | RA | RB | XO | 1 |
| PO | FRTp |  | S III | FRBp | XO | RC |
| PO | FRTp |  | SP III | FRBp | XO | Rc |
| PO | RS |  | III | RB | XO | 1 |
| PO | RS |  | 1 RIC PR R | RB | XO | 1 |
| PO | RS |  | III L | III | XO | / |
| PO | RS |  | 1 SR | III | XO | / |
| PO | RS |  | BFA // | III | XO | 1 |
| PO | RS |  | RA | III | XO | / |
| PO | RS |  | RA | I/I | XO | 1 |
| PO | RS |  | RA | III | XO | RC |
| PO | RS |  | RA | FC | XO | / |
| PO | RS |  | RA | NB | XO | 1 |
| PO | RS |  | RA | SH | XO | Rc |
| PO | RS |  | RA | RB | XO | 1 |
| PO | RS |  | RA | RB | XO | 1 |
| PO | RS |  | RA | RB | XO | Rc |
| PO | RSp |  | RA | RB | XO | 1 |
| PO | RT |  | III | III | XO | / |
| PO | RT |  | III | RB | XO | 1 |

Figure 17. X instruction format

## Version 3.0

| 7891011121314151617181920212223242526 |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| PO | RT | III | RB | XO | 1 |
| PO | RT | III L | III | XO | 1 |
| PO | RT | / 1 SR | III | XO | 1 |
| PO | RT | RA | FC | XO | 1 |
| PO | RT | RA | NB | XO | 1 |
| PO | RT | RA | RB | XO | 1 |
| PO | RT | RA | RB | XO | EH |
| PO | RTp | RA | RB | XO | EH |
| PO | S | RA | III | XO | SX |
| PO | S | RA | RB | XO | SX |
| PO | T | RA | III | XO | TX |
| PO | T | RA | RB | XO | TX |
| PO | T | EO | M8 | XO | TX |
| PO | TH | RA | RB | XO | 1 |
| PO | TO | RA | SI | XO | 1 |
| PO | TO | RA | RB | XO | 1 |
| PO | TO | RA | RB | XO | 1 |
| PO | VRS | RA | RB | XO | 1 |
| PO | VRT | RA | RB | XO | 1 |
| PO | VRT | VRA | VRB | XO | 1 |
| PO | VRT | VRA | VRB | XO | RO |
| PO | VRT | XO | VRB | XO | 1 |
| PO | VRT | XO | VRB | XO | RO |

Figure 17. X instruction format

### 1.6.16 XFL-FORM

| 67 |  |  | 1516 | 21 |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| PO | L | FLM | M | FRB | XO | RC |

Figure 18. XFL instruction format

### 1.6.17 XFX-FORM

|  | 6 | 1112 | 1516 | 2021 |  | 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | III |  | 1 |  | XO | 1 |
| PO | RS | 0 | FXM | 1 | XO | 1 |
| PO | RS | 1 | FXM | 1 | XO | 1 |
| PO | RS |  | spr |  | XO | / |
| PO | RT |  | BHRBE |  | XO | 1 |
| PO | RT | 0 | III | 1 | XO | 1 |
| PO | RT | 1 | FXM | 1 | XO | 1 |
| PO | RT |  | spr |  | XO | 1 |
| PO | RT |  | tbr |  | XO | 1 |

Figure 19. XFX instruction format

### 1.6.18 XL-FORM

|  | 69 | 1114 | 61920 |  | 31 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| PO | III | III | III | xO | 1 |
| PO | III | III | III S | XO | 1 |
| PO | BF I/ | BFA // | III | XO | 1 |
| PO | BO | BI | III ${ }^{\text {BH }}$ | XO | LK |
| PO | BT | BA | BB | XO | 1 |

Figure 20. XL instruction format

### 1.6.19 XO-FORM



Figure 21. XO instruction format

### 1.6.20 XS-FORM



Figure 22. XS instruction format

### 1.6.21 XX2-FORM



Figure 23. XX2 instruction format

### 1.6.22 XX3-FORM

|  | 6 | 9 | 16 |  | 2122 |  | 24 | $\begin{aligned} & 293031 \\ & \|\operatorname{AX}\| \operatorname{BX}\|/\| \end{aligned}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | BF | // | A | B |  |  |  |  |
| PO | T |  | A | B | 0 | DM | XO | AX BX TX |
| PO | T |  | A | B | 0 | SHW | XO | AX BX TX |
| PO | T |  | A | B | RC |  | O | AX BX TX |
| PO | T |  | A | B |  |  |  | AX BX TX |

Figure 24. XX3 instruction format

### 1.6.23 XX4-FORM



Figure 25. XX4 instruction format

### 1.6.24 Z22-FORM

| 6 |  |  |  |  |  | 11 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | BF | $/ /$ | FRA | DCM | XO | $/$ |
| PO | BF | $/ /$ | FRA | DGM | XO | $/$ |
| PO | BF | $/ /$ | FRAp | DCM | XO | $/$ |
| PO | BF | $/ /$ | FRAp | DGM | XO | $/$ |
| PO | FRT | FRA | SH | XO | Rc |  |
| PO | FRTp | FRAp | SH | XO | Rc |  |

Figure 26. Z22 instruction format

### 1.6.25 Z23-FORM

| 0 | 6 | 1516 |  | $21 \quad 23$ |  | 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| PO | FRT | III | FRB | RMC | XO | Rc |
| PO | FRT | FRA | FRB | RMC | XO | Rc |
| PO | FRT | TE | FRB | RMC | XO | Rc |
| PO | FRTp | I/I | FRBp | RMC | XO | Rc |
| PO | FRTp | FRA | FRBp | RMC | XO | Rc |
| PO | FRTp | FRAp | FRBp | RMC | XO | Rc |
| PO | FRTp | TE | FRBp | RMC | XO | Rc |
| PO | VRT | III | VRB | RMC | XO | / |
| PO | VRT | III R | VRB | RMC | XO | EX |

Figure 27. Z23 instruction format

### 1.7 Instruction Fields

## A (6)

Field used by the tbegin. instruction to specify an implementation-specific function.

Field used by the tend. instruction to specify the completion of the outer transaction and all nested transactions.

Formats: X
AA (30)
Absolute Address.
0 The immediate field represents an address relative to the current instruction address. For I-form branches the effective address of the branch target is the sum of the LI field sign-extended to 64 bits and the address of the branch instruction. For B-form branches the effective address of the branch target is the sum of the BD field sign-extended to 64 bits and the address of the branch instruction.

1 The immediate field represents an absolute address. For I-form branches the effective address of the branch target is the LI field sign-extended to 64 bits. For B-form branches the effective address of the branch target is the BD field sign-extended to 64 bits.

Formats: B, I

## AX,A (29,11:15)

Fields that are concatenated to specify a VSR to be used as a source.

Formats: XX3, XX4

## BA (11:15)

Field used to specify a bit in the CR to be used as a source.

Formats: XL

## BB (16:20)

Field used to specify a bit in the CR to be used as a source.

Formats: XL
BC (21:25)
Field used to specify a bit in the CR to be used as a source.

Formats: A
BD (16:29)
Immediate field used to specify a 14-bit signed two's complement branch displacement which is concatenated on the right with 0 bOO and sign-extended to 64 bits.
Formats: B
BF (6:8)
Field used to specify one of the CR fields or one of the FPSCR fields to be used as a target.
Formats: D, X, XL, XX2, XX3, Z22
BFA (11:13)
Field used to specify one of the CR fields or one of the FPSCR fields to be used as a source.

Formats: X, XL

## BH (19:20)

Field used to specify a hint in the Branch Conditional to Link Register and Branch Conditional to Count Register instructions. The encoding is described in Section 2.5, "Branch Instructions".
Formats: XL

## BHRB (11:20)

Field used to identify the BHRB entry to be used as a source by the Move From Branch History Rolling Buffer instruction.

Formats: X

## BI (11:15)

Field used to specify a bit in the CR to be tested by a Branch Conditional instruction.

Formats: B, XL

## BO (6:10)

Field used to specify options for the Branch Conditional instructions. The encoding is described in Section 2.5, "Branch Instructions".
Formats: $\mathrm{B}, \mathrm{XL}, \mathrm{X}, \mathrm{XL}$

## BT (6:10)

Field used to specify a bit in the CR or in the FPSCR to be used as a target.

Formats: XL

## BX,B (30,16:20)

Fields that are concatenated to specify a VSR to be used as a source.

Formats: XX2, XX3, XX4
CT (7:10)
Field used in X-form instructions to specify a cache target (see Section 4.3.2 of Book II).

Formats: X
CX,C (28,21:25)
Fields that are concatenated to specify a VSR to be used as a source.

Formats: XX4
D (16:31)
Immediate field used to specify a 16 -bit signed two's complement integer which is sign-extended to 64 bits.

Formats: D
d0,d1,d2 (16:25,11:15,31)
Immediate fields that are concatenated to specify a 16-bit signed two's complement integer which is sign-extended to 64 bits.

Formats: DX
dc,dm,dx (25,29,11:15)
Immediate fields that are concatenated to specify Data Class Mask.

Formats: XX2
DCM (16:21)
Immediate field used to specify Data Class Mask.
Formats: Z22
DCMX (9:15)
Immediate field used to specify Data Class Mask.
Formats: X, XX2
DGM (16:21)
Immediate field used as the Data Group Mask.
Formats: Z22

## DM (22:23)

Immediate field used by xxpermdi instruction as doubleword permute control.

Formats: XX3

## DQ (16:27)

Immediate field used to specify a 12-bit signed two's complement integer which is concatenated on the right with $0 b 0000$ and sign-extended to 64 bits.

Formats: DQ

## DS (16:29)

Immediate field used to specify a 14-bit signed two's complement integer which is concatenated on the right with $0 b 00$ and sign-extended to 64 bits.

Formats: DS
EH (31)
Field used to specify a hint in the Load and Reserve instructions. The meaning is described in Section 4.6.2, "Load and Reserve and Store Conditional Instructions", in Book II.

Formats: X

## EO (11:12)

Expanded opcode field
Formats: XX1
EO (11:15)
Expanded opcode field
Formats: VX, X, XX2
EX (31)
Field used to specify Inexact form of round to quad-precision integer.
Formats: X
FC (16:20)
Field used to specify load/store atomic function code.

Formats: X

## FLM (7:14)

Field mask used to identify the FPSCR fields that are to be updated by the mtfsf instruction.

Formats: XFL
FRA (11:15)
Field used to specify a FPR to be used as a source.

Formats: A, X, Z22, Z23

## FRAp (11:15)

Field used to specify an even/odd pair of FPRs to be concatenated and used as a source.

Formats: X, Z22, Z23
FRB (16:20)
Field used to specify an FPR to be used as a source.

Formats: A, X, XFL, Z23
FRBp (16:20)
Field used to specify an even/odd pair of FPRs to be concatenated and used as a source.

Formats: X, Z23

## FRC (21:25)

Field used to specify an FPR to be used as a source.

Formats: A

## FRS (6:10)

Field used to specify an FPR to be used as a source.

Formats: D, X

## FRSp (6:10)

Field used to specify an even/odd pair of FPRs to be concatenated and used as a source.

Formats: DS, X

## FRT (6:10)

Field used to specify an FPR to be used as a target.

Formats: A, D, X, Z22, Z23
FRTp (6:10)
Field used to specify an even/odd pair of FPRs to be concatenated and used as a target.
Formats: DS, X, Z22, Z23

## FXM (12:19)

Field mask used to identify the CR fields that are to be written by the mtcrf and mtocrf instructions, or read by the mfocrf instruction.

Formats: XFX

## IB (16:20)

Immediate field used to specify a 5-bit signed integer.

Formats: MDS

## IH (8:10)

Field used to specify a hint in the SLB Invalidate All instruction. The meaning is described in Section 5.9.3.2, "SLB Management Instructions", in Book III.

Formats: X

## IMM8 (13:20)

Immediate field used to specify a 8-bit integer.
Formats: XX1
IS (6:10)
Immediate field used to specify a 5-bit signed integer.
Formats: MDS

## L(6)

Field used to specify whether the mtfsf instruction updates the entire FPSCR.

Formats: XFL

## L (9:10)

Field used by the Data Cache Block Flush instruction (see Section 4.3.2 of Book II) and also by the Synchronize instruction (see Section 4.6.3 of Book II).

Formats: X
L (10)
Field used to specify whether a fixed-point Compare instruction is to compare 64-bit numbers or 32-bit numbers.
Field used by the Copy and Paste instructions to indicate the first Copy and the last Paste instructions in the move group.
Formats: D, X
L (15)
Field used by the Move To Machine State Register instruction (see Book III).
Field used by the SLB Move From Entry VSID and SLB Move From Entry ESID instructions for imple-mentation-specific purposes.
Formats: X
L (14:15)
Field used by the Deliver A Random Number instruction (see Section 3.3.9, "Fixed-Point Arithmetic Instructions") to choose the random number format.

Formats: X
LEV (20:26)
Field used by the System Call instructions.
Formats: SC
LI (6:29)
Immediate field used to specify a 24 -bit signed two's complement integer which is concatenated on the right with 0 bOO and sign-extended to 64 bits.

Formats: I
LK (31)
LINK bit.
0 Do not set the Link Register.
1 Set the Link Register. The address of the instruction following the Branch instruction is placed into the Link Register.

Formats: B, I, XL
MB (21:25)
Field used in M-form instructions to specify the first 1-bit of a 64-bit mask, as described in Section 3.3.14, "Fixed-Point Rotate and Shift Instructions" on page 100.
Formats: M

## 18 Power ISA ${ }^{\text {TM }}$ - Book I

## mb (21:26)

Field used in MD-form and MDS-form instructions to specify the first 1-bit of a 64-bit mask, as described in Section 3.3.14, "Fixed-Point Rotate and Shift Instructions" on page 100.

Formats: MD, MDS
me (21:26)
Field used in MD-form and MDS-form instructions to specify the last 1-bit of a 64-bit mask, as described in Section 3.3.14, "Fixed-Point Rotate and Shift Instructions" on page 100.

Formats: MD, MDS

## ME (26:30)

Field used in M-form instructions to specify the last 1-bit of a 64-bit mask, as described in Section 3.3.14, "Fixed-Point Rotate and Shift Instructions" on page 100.

Formats: M
NB (16:20)
Field used to specify the number of bytes to move in an immediate Move Assist instruction.

Formats: X
OE (21)
Field used by XO-form instructions to enable setting $O V$ and SO in the XER.

Formats: XO
PO (0:5)
Primary opcode.
Formats: all
PRS (14)
Field used to specify whether to invalidate pro-cess- or partition-scoped entries for tlbie[I].

Formats: X
PS (22)
Field used to specify preferred sign for BCD operations.

Formats: VX
PT (28:31)
Immediate field used to specify a 4-bit unsigned value.

Formats: DQ
R (10)
Field used by the tbegin. instruction to specify the start of a ROT.

Formats: X

R (15)
Immediate field that specifies whether the RMC is specifying the primary or secondary encoding

Field used to specify whether to invalidate Radix Tree or HPT entries for tlbie[ [] .

Formats: X, Z23
RA (11:15)
Field used to specify a GPR to be used as a source or as a target.

Formats: A, D, DQ, DQE, DS, M, MD, MDS, TX, VA, VX, X, XO, XS, XX1

RB (16:20)
Field used to specify a GPR to be used as a source.

Formats: A, M, MDS, VA, X, XO, XX1
Rc (21)
RECORD bit.
0 Do not alter the Condition Register.
1 Set Condition Register Field 6 as described in Section 2.3.1, "Condition Register" on page 30 .

Formats: VC, XX3
RC (21:25)
Field used to specify a GPR to be used as a source.

Formats: VA
Rc (31)
RECORD bit.
0 Do not alter the Condition Register.
1 Set Condition Register Field 0 or Field 1 as described in Section 2.3.1, "Condition Register" on page 30.

Formats: A, M, MD, MDS, X, XFL, XO, XS, Z22, Z23
RIC (12:13)
Field used to specify what types of entries to invalidate for tlbie[I].
Formats: X
RMC (21:22)
Immediate field used for DFP rounding mode control.

Formats: Z23
RO (31)
Round to Odd override
Formats: X

## RS (6:10)

Field used to specify a GPR to be used as a source.

Formats: D, DS, M, MD, MDS, X, XFX, XS

## RSp (6:10)

Field used to specify an even/odd pair of GPRs to be concatenated and used as a source.

Formats: DS, X

## RT (6:10)

Field used to specify a GPR to be used as a target.
Formats: A, D, DQE, DS, DX, VA, VX, X, XFX, XO, XX2

## RTp (6:10)

Field used to specify an even/odd pair of GPRs to be concatenated and used as a target.

Formats: DQ, X

## S (11)

Immediate field that specifies signed versus unsigned conversion.

Formats: X
S (20)
Immediate field that specifies whether or not the rfebb instruction re-enables event-based branches.

Formats: XL

## SH (16:20)

Field used to specify a shift amount.
Formats: $\mathrm{M}, \mathrm{X}$

## SH (16:21)

Field used to specify a shift amount.
Formats: Z22
sh (30,16:20)
Fields that are concatenated to specify a shift amount.

Formats: MD, XS

## SHB (22:25)

Field used to specify a shift amount in bytes.
Formats: VA

## SHW (22:23)

Field used to specify a shift amount in words.
Formats: XX3

## SI (16:20)

Immediate field used to specify a 5-bit signed integer.
Formats: X

## SI (16:31)

Immediate field used to specify a 16 -bit signed integer.
Formats: D

## SIM (11:15)

Immediate field used to specify a 5-bit signed integer.
Formats: VX
SP (11:12)
Immediate field that specifies signed versus unsigned conversion.

Formats: X
SPR (11:20)
Field used to specify a Special Purpose Register for the mtspr and mfspr instructions.

Formats: X

## SR (12:15)

Field used by the Segment Register Manipulation instructions (see Book III).
Formats: X

## SX,S (28,6:10)

Fields SX and S are concatenated to specify a VSR to be used as a source.
Formats: DQ

## SX,S (31,6:10)

Fields SX and S are concatenated to specify a VSR to be used as a source.

Formats: XX1
TBR (11:20)
Field used by the Move From Time Base instruction (see Section 6.1 of Book II).
Formats: X
TE (11:15)
Immediate field that specifies a DFP exponent.
Formats: Z23
TH (6:10)
Field used by the data stream variant of the dcbt and dcbtst instructions (see Section 4.3.2 of Book II).

Formats: X
TO (6:10)
Field used to specify the conditions on which to trap. The encoding is described in Section 3.3.10.1, "Character-Type Compare Instructions" on page 87.
Formats: TX, X

## 20

## TX,T (28,6:10)

Fields that are concatenated to specify a VSR to be used as either a target.

Formats: DQ
TX,T (31,6:10)
Fields that are concatenated to specify a VSR to be used as either a target or a source.

Formats: XX1, XX2, XX3, XX4
U (16:19)
Immediate field used as the data to be placed into a field in the FPSCR.

Formats: X
UI (16:20)
Immediate field used to specify a 5-bit unsigned integer.

Formats: TX
UI (16:31)
Immediate field used to specify a 16-bit unsigned integer.

Formats: D
UIM (11:15)
Immediate field used to specify a 5-bit unsigned integer.
I Formats: $\mathrm{VX}, \mathrm{X}$
UIM (12:15)
Immediate field used to specify a 4-bit unsigned integer.
| Formats: VX, XX2
UIM (13:15)
Immediate field used to specify a 3-bit unsigned integer.

Formats: VX
UIM (14:15)
Immediate field used to specify a 2-bit unsigned integer.

Formats: VX, XX2

## VRA (11:15)

Field used to specify a VR to be used as a source.
Formats: VA, VC, VX
VRB (16:20)
Field used to specify a VR to be used as a source.
Formats: VA, VC, VX
VRC (21:25)
Field used to specify a VR to be used as a source.
Formats: VA

VRS (6:10)
Field used to specify a VR to be used as a source.
| Formats: DS, X
VRT (6:10)
Field used to specify a VR to be used as a target.
I Formats: DS, VA, VC, VX, X
W (15)
Field used by the mtfsfi and mtfsf instructions to specify the target word in the FPSCR.
Formats: X, XFL
I WC (9:10)
Field used to specify the condition or conditions that cause instruction execution to resume after executing a wait instruction (see Section 4.6.4 of Book II).

Formats: X
XBI (21:24)
Field used to specify a bit in the XER.
Formats: MDS, MDS, TX

## XO (21,23:31)

Extended opcode field.
Formats: VX
XO (21:24,26:28)
Extended opcode field.
Formats: XX2
XO (21:24:28)
Extended opcode field.
Formats: XX3
XO (21:28)
Extended opcode field.
Formats: XX3
XO (21:29)
Extended opcode field.
Formats: XS, XX2
XO (21:30)
Extended opcode field.
Formats: X, XFL, XFX, XL, XX1
XO (21:31)
Extended opcode field.
Formats: VX
XO (22:30)
Extended opcode field.
Formats: XO, XX3, Z22

## XO (22:31)

Extended opcode field.
Formats: VC
XO (23:30)
Extended opcode field.
Formats: X, Z23
XO (25:30)
Extended opcode field.
Formats: TX

## XO (26:27)

Extended opcode field.
Formats: XX4

## XO (26:30)

Extended opcode field.
Formats: A, DX
XO (26:31)
Extended opcode field.
Formats: VA
XO (27:29)
Extended opcode field.
Formats: MD

## XO (27:30)

Extended opcode field.
Formats: MDS

## XO (29:31)

Extended opcode field.
Formats: DQ
XO (30)
Extended opcode field.
Formats: SC
XO (30:31)
Extended opcode field.
Formats: DQE, DS, SC

### 1.8 Classes of Instructions

An instruction falls into exactly one of the following three classes:

Defined
Illegal
Reserved
The class is determined by examining the opcode, and the extended opcode if any. If the opcode, or combination of opcode and extended opcode, is not that of a defined instruction or a reserved instruction, the instruction is illegal.

### 1.8.1 Defined Instruction Class

This class of instructions contains all the instructions defined in this document.

A defined instruction can have preferred and/or invalid forms, as described in Section 1.9.1, "Preferred Instruction Forms" and Section 1.9.2, "Invalid Instruction Forms".

### 1.8.2 Illegal Instruction Class

This class of instructions contains the set of instructions described in Appendix A of Book Appendices. Illegal instructions are available for future extensions of the Power ISA ; that is, some future version of the Power ISA may define any of these instructions to perform new functions.

Any attempt to execute an illegal instruction will cause the system illegal instruction error handler to be invoked and will have no other effect.

An instruction consisting entirely of binary $0 s$ is guaranteed always to be an illegal instruction. This increases the probability that an attempt to execute data or uninitialized storage will result in the invocation of the system illegal instruction error handler.

### 1.8.3 Reserved Instruction Class

This class of instructions contains the set of instructions described in Appendix B of Book Appendices.
Reserved instructions are allocated to specific purposes that are outside the scope of the Power ISA.
Any attempt to execute a reserved instruction will:
■ perform the actions described by the implementation if the instruction is implemented; or

- cause the system illegal instruction error handler to be invoked if the instruction is not implemented.


## 22

### 1.9 Forms of Defined Instructions

### 1.9.1 Preferred Instruction Forms

Some of the defined instructions have preferred forms. For such an instruction, the preferred form will execute in an efficient manner, but any other form may take significantly longer to execute than the preferred form.

Instructions having preferred forms are:

- the Condition Register Logical instructions
- the Load Quadword instruction
- the Move Assist instructions
- the Or Immediate instruction (preferred form of no-op)
- the Move To Condition Register Fields instruction


### 1.9.2 Invalid Instruction Forms

Some of the defined instructions can be coded in a form that is invalid. An instruction form is invalid if one or more fields of the instruction, excluding the opcode field(s), are coded incorrectly in a manner that can be deduced by examining only the instruction encoding.

In general, any attempt to execute an invalid form of an instruction will either cause the system illegal instruction error handler to be invoked or yield boundedly undefined results. Exceptions to this rule are stated in the instruction descriptions.

Some instruction forms are invalid because the instruction contains a reserved value in a defined field (see Section 1.3.3 on page 5); these invalid forms are not discussed further. All other invalid forms are identified in the instruction descriptions.

References to instructions elsewhere in this document assume the instruction form is not invalid, unless otherwise stated or obvious from context.

## Assembler Note

Assemblers should report uses of invalid instruction forms as errors.

### 1.9.3 Reserved-no-op Instructions

Reserved-no-op instructions include the following extended opcodes under primary opcode 31: 530, 562, 594, 626, 658, 690, 722, and 754.
Reserved-no-op instructions are provided in the architecture to anticipate the eventual adoption of performance hint instructions to the architecture. For these instructions, which cause no visible change to architected state, employing a reserved-no-op opcode will allow software to use this new capability on new implementations that support it while remaining compatible
with existing implementations that may not support the new function.

When a reserved-no-op instruction is executed, no operation is performed.

Reserved-no-op instructions are not assigned instruction names or mnemonics. There are no individual descriptions of reserved-no-op instructions in this document.

### 1.10 Exceptions

There are two kinds of exception, those caused directly by the execution of an instruction and those caused by an asynchronous event. In either case, the exception may cause one of several components of the system software to be invoked.

The exceptions that can be caused directly by the execution of an instruction include the following:
■ an attempt to execute an illegal instruction, or an attempt by an application program to execute a "privileged" instruction (see Book III) (system illegal instruction error handler or system privileged instruction error handler)

■ the execution of a defined instruction using an invalid form (system illegal instruction error handler or system privileged instruction error handler)

- an attempt to execute an instruction that is not provided by the implementation (system illegal instruction error handler)
- an attempt to access a storage location that is unavailable (system instruction storage error handler or system data storage error handler)
■ an attempt to access storage with an effective address alignment that is invalid for the instruction (system alignment error handler)
■ the execution of a System Call or System Call Vectored instruction (system service program)
■ the execution of a Trap instruction that traps (system trap handler)
- the execution of a floating-point instruction that causes a floating-point enabled exception to exist (system floating-point enabled exception error handler)
- the execution of an auxiliary processor instruction that causes an auxiliary processor enabled exception to exist (system auxiliary processor enabled exception error handler)

The exceptions that can be caused by an asynchronous event are described in Book III.
The invocation of the system error handler is precise, except that the invocation of the auxiliary processor enabled exception error handler may be imprecise, and
if one of the imprecise modes for invoking the system floating-point enabled exception error handler is in effect (see page 133), then the invocation of the system floating-point enabled exception error handler may also be imprecise. When the system error handler is invoked imprecisely, the excepting instruction does not appear to complete before the next instruction starts (because one of the effects of the excepting instruction, namely the invocation of the system error handler, has not yet occurred).
Additional information about exception handling can be found in Book III.

### 1.11 Storage Addressing

A program references storage using the effective address computed by the processor when it executes a Storage Access or Branch instruction (or certain other instructions described in Book II and Book III), or when it fetches the next sequential instruction.

Bytes in storage are numbered consecutively starting with 0 . Each number is the address of the corresponding byte.
The byte ordering (Big-Endian or Little-Endian) for a storage access is specified by the operating system. This byte ordering is also referred to as the Endian mode and it applies to both data accesses and instruction fetches. The Endian mode is specified by the LE mode bit (see Section 3.2.1 of Book III), which applies to all of storage.

### 1.11.1 Storage Operands

A storage operand may be a byte, a halfword, a word, a doubleword, or a quadword, or, for the Load/Store Multiple and Move Assist instructions, a sequence of bytes (Move Assist) or words (Load/Store Multiple). The address of a storage operand is the address of its first byte (i.e., of its lowest-numbered byte). An instruction for which the storage operand is a byte is said to cause a byte access, and similarly for halfword, word, doubleword, and quadword.

The length of the storage operand is the number of bytes (of the storage operand) that the instruction would access in the absence of invocations of the system error handler. The length is generally implied by the name of the instruction (equivalently, by the opcode, and extended opcode if any). For example, the length of the storage operand of a Load Word and Zero, Load Floating-Point Single, and Load Vector Element Word instruction is four bytes (one word), and the length of a Store Quadword, Store Floating-Point Double Pair, and Store VSX Vector Word*4 instruction is 16 bytes (one quadword). The only exceptions are the Load/Store Multiple and Move Assist instructions, for which the length of the storage operand is implied by the identity of the specified source or target register
(Load/Store Multiple), or by an immediate field in the instruction or the contents of a field in the XER (Move Assist), as well as by the name of the instruction. For example, the length of the storage operand of a Load Multiple Word instruction for which the specified target register is GPR 20 is 48 bytes ((32-20)x4), and the length of the storage operand of a Load String Word Immediate instruction for which the immediate field contains the number 20 is 20 bytes.

The storage operand of a Load or Store instruction other than a Load/Store Multiple or Move Assist instruction is said to be aligned if the address of the storage operand is an integral multiple of the storage operand length; otherwise it is said to be unaligned. See the following table. (The storage operand of a Load/Store Multiple or Move Assist instruction is neither said to be aligned nor said to be unaligned. Its alignment properties are described, when necessary, using terms such as "word-aligned", which are defined below.)

| Operand | Length | Addr $_{60: 63}$ if aligned |
| :--- | :--- | :--- |
| Byte | 8 bits | $x \times x \times$ |
| Halfword | 2 bytes | $x x x 0$ |
| Word | 4 bytes | $x x 00$ |
| Doubleword | 8 bytes | $x 000$ |
| Quadword | 16 bytes | 0000 |
| Nual | An" in |  |

Note: An " $x$ " in an address bit position indicates that the bit can be 0 or 1 independent of the contents of other bits in the address.

The concept of alignment is also applied more generally, to any datum in storage.

- A datum having length that is an integral power of 2 is said to be aligned if its address is an integral multiple of its length.
- A datum of any length is said to be half-word-aligned (or aligned at a halfword boundary) if its address is an integral multiple of 2, word-aligned (or aligned at a word boundary) if its address is an integral multiple of 4 , etc. (All data in storage is byte-aligned.)

The concept of alignment can also be applied to data in registers, with the "address" of the datum interpreted as the byte number of the datum in the register. E.g., a word element (4 bytes) in a Vector Register is said to be aligned if its byte number is an integral multiple of 4 .

## Programming Note

The technical literature sometimes uses the term "naturally aligned" to mean "aligned."

Versions of the architecture that precede Version 2.07 also used "naturally aligned" as defined above. The term was dropped from the architecture in Version 2.07 because it seemed to mean different things to different readers and is not needed.

Some instructions require their storage operands to have certain alignments. In addition, alignment may affect performance. In general, the best performance is obtained when storage operands are aligned.

When a storage operand of length N bytes starting at effective address EA is copied between storage and a register that is R bytes long (i.e., the register contains bytes numbered from 0 , most significant, through R-1, least significant), the bytes of the operand are placed into the register or into storage in a manner that depends on the byte ordering for the storage access as shown in Figure 28, unless otherwise specified in the instruction description.

| Big-Endian Byte Ordering |  |
| :--- | :--- |
| Load | Store |
| for $i=0$ to $N-1:$ | for $i=0$ to $N-1:$ |
| $R T_{(R-N)+i} \leftarrow M E M(E A+i, 1)$ | MEM $(E A+i, 1) \leftarrow(R S)_{(R-N)+i}$ |
| Little-Endian Byte Ordering |  |
| Load | Store |
| for $i=0$ to $N-1:$ | for $i=0$ to $N-1:$ |
| $R T_{(R-1)-i} \leftarrow M E M(E A+i, 1)$ | MEM $(E A+i, 1) \leftarrow(R S)_{(R-1)-i}$ |

## Notes:

1. In this table, subscripts refer to bytes in a register rather than to bits as defined in Section 1.3.2.
2. This table does not apply to the Ivebx, Ivehx, Ivewx, stvebx, stvehx, and stvewx instructions.

Figure 29 shows an example of a C language structure s containing an assortment of scalars and one character string. The value assumed to be in each structure element is shown in hex in the C comments; these values are used below to show how the bytes making up each structure element are mapped into storage. It is assumed that structure s is compiled for 32-bit mode or for a 32-bit implementation. (This affects the length of the pointer to c.)
C structure mapping rules permit the use of padding (skipped bytes) in order to align the scalars on desirable boundaries. Figures 30 and 31 show each scalar as aligned. This alignment introduces padding of four bytes between $\mathbf{a}$ and $\mathbf{b}$, one byte between $\mathbf{d}$ and $\mathbf{e}$, and two bytes between e and $\mathbf{f}$. The same amount of padding is present for both Big-Endian and Little-Endian mappings.
The Big-Endian mapping of structure $\mathbf{s}$ is shown in Figure 30. Addresses are shown in hex at the left of each doubleword, and in small figures below each byte. The contents of each byte, as indicated in the $C$ example in Figure 29, are shown in hex (as characters for the elements of the string).

The Little-Endian mapping of structure s is shown in Figure 31. Doublewords are shown laid out from right to left, which is the common way of showing storage maps for processors that implement only Little-Endian byte ordering.

Figure 28. Storage operands and byte ordering

| struct \{ int | a; | /* 0x1112_1314 | word |
| :---: | :---: | :---: | :---: |
| double | b; | /* 0x2122_2324_2526_2728 | doubleword |
| char * | c; | /* 0x3132_3334 | word |
| char | d[7]; | /* 'A', 'B', 'C', 'D', 'E', 'F', 'G' | array of bytes |
| short | e; | /* 0x5152 | halfword |
| int | f; | /* 0x6162_6364 | word |

Figure 29. $C$ structure ' $s$ ', showing values of elements

| 00 | $\begin{gathered} 11 \\ 00 \end{gathered}$ | $\begin{gathered} 12 \\ 01 \end{gathered}$ | $\begin{aligned} & 13 \\ & 02 \end{aligned}$ | $\begin{gathered} 14 \\ 03 \end{gathered}$ | 04 | 05 | 06 | 07 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 08 | $\begin{array}{r} 21 \\ 08 \\ \hline \end{array}$ | $\begin{gathered} 22 \\ 09 \end{gathered}$ | $\begin{array}{r} 23 \\ 0 A \\ \hline \end{array}$ | $\begin{array}{r} 24 \\ \text { OB } \\ \hline \end{array}$ | 25 00 | $\begin{gathered} 26 \\ 00 \end{gathered}$ | $\begin{gathered} 27 \\ 0 \mathrm{OE} \\ \hline \end{gathered}$ | $\begin{gathered} 28 \\ 0 \\ \hline \end{gathered}$ |
| 10 | $\begin{aligned} & 31 \\ & 10 \\ & \hline \end{aligned}$ | $\begin{gathered} 32 \\ 11 \\ \hline \end{gathered}$ | $\begin{aligned} & 33 \\ & 12 \\ & \hline \end{aligned}$ |  |  | $\begin{gathered} \hline \quad \mathrm{B}^{\prime} \\ 15 \\ \hline \end{gathered}$ | $\begin{aligned} & \text { 'C' } \\ & 16 \\ & \hline \end{aligned}$ | $17$ |
| 18 | $\begin{aligned} & {f2b129f28-79aa-49ca-9622-3b705b5444ef}^{\prime} F^{\prime} \\ & 19 \\ & \hline \end{aligned}$ | $\begin{aligned} & { }^{\prime} \mathrm{G}^{\prime} \\ & 1 \mathrm{~A} \end{aligned}$ | 1B |  | $\begin{array}{r} 52 \\ 10 \\ \hline \end{array}$ | 1 E | 1 F |  |
| 20 | $\begin{array}{r}61 \\ 20 \\ \hline\end{array}$ | $\begin{array}{r} 62 \\ 21 \\ \hline \end{array}$ | $\begin{aligned} & 63 \\ & 22 \\ & \hline \end{aligned}$ | $\begin{array}{r} \hline 64 \\ 23 \\ \hline \end{array}$ |  |  |  |  |



Figure 31. Little-Endian mapping of structure ' $s$ '
Figure 30. Big-Endian mapping of structure ' s '

### 1.11.2 Instruction Fetches

Instructions are always four bytes long and word-aligned.

When an instruction starting at effective address EA is fetched from storage, the relative order of the bytes within the instruction depend on the byte ordering for the storage access as shown in Figure 32.


Figure 32. Instructions and byte ordering
Figure 33 shows an example of a small assembly language program $\mathbf{p}$.

| loop: |  |  |
| :--- | :--- | :--- |
|  | cmplwi | r5,0 |
|  | beq | done |
|  | lwzux | r4,r5,r6 |
|  | add | r7,r7,r4 |
|  | subi | r5,r5,4 |
|  | b | loop |
| done: |  |  |
|  | stw | $r 7$, total |

Figure 33. Assembly language program ' $\mathbf{p}$ '
The Big-Endian mapping of program $\mathbf{p}$ is shown in Figure 34 (assuming the program starts at address 0).

| 00 | loop: cmplwi r5,0 |  |  |  | beq done |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 |
| 08 | lwzux r4,r5,r6 |  |  |  | add r7,r7,r4 |  |  |  |
|  | 08 | 09 | OA | OB | OC | OD | OE | OF |
| 10 | subi r5,r5,4 |  |  |  | b loop |  |  |  |
|  | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 |
| 18 | done: stw r7,total |  |  |  |  |  |  |  |
|  | 18 | 19 | 1A | 1B | 1 C | 1D | 1 E | 1F |

Figure 34. Big-Endian mapping of program ' $p$ '
The Little-Endian mapping of program $\mathbf{p}$ is shown in Figure 35.

| beq done |  |  |  | loop: cmplwi r5,0 |  |  |  | 00 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 07 | 06 | 05 | 04 | 03 | 02 | 01 | 00 |  |
| add r7, r7, r4 |  |  |  | 1wzux r4,r5,r6 |  |  |  | 08 |
| OF | OE | OD | 0C | OB | OA | 09 | 08 |  |
| b loop |  |  |  | subi r5,r5,4 |  |  |  | 10 |
| 17 | 16 | 15 | 14 | 13 | 12 | 11 | 10 |  |
|  |  |  |  | done | st | r7, | tal | 18 |
| 1F | 1E | 1D | 1C | 1B | 1A | 19 | 18 |  |

Figure 35. Little-Endian mapping of program ' $p$ '

The terms Big-Endian and Little-Endian come from Part I, Chapter 4, of Jonathan Swift's Gulliver's Travels. Here is the complete passage, from the edition printed in 1734 by George Faulkner in Dublin.
... our Histories of six Thousand Moons make no Mention of any other Regions, than the two great Empires of Lilliput and Blefuscu. Which two mighty Powers have, as I was going to tell you, been engaged in a most obstinate War for six and thirty Moons past. It began upon the following Occasion. It is allowed on all Hands, that the primitive Way of breaking Eggs before we eat them, was upon the larger End: But his present Majesty's Grand-father, while he was a Boy, going to eat an Egg, and breaking it according to the ancient Practice, happened to cut one of his Fingers. Whereupon the Emperor his Father, published an Edict, commanding all his Subjects, upon great Penalties, to break the smaller End of their Eggs. The People so highly resented this Law, that our Histories tell us, there have been six Rebellions raised on that Account; wherein one Emperor lost his Life, and another his Crown. These civil Commotions were constantly fomented by the Monarchs of Blefuscu; and when they were quelled, the Exiles always fled for Refuge to that Empire. It is computed that eleven Thousand Persons have, at several Times, suffered Death, rather than submit to break their Eggs at the smaller End. Many hundred large Volumes have been published upon this Controversy: But the Books of the Big-Endians have been long
forbidden, and the whole Party rendered incapable by Law of holding Employments. During the Course of these Troubles, the Emperors of Blefuscu did frequently expostulate by their Ambassadors, accusing us of making a Schism in Religion, by offending against a fundamental Doctrine of our great Prophet Lustrog, in the fifty-fourth Chapter of the Brundrecal, (which is their Alcoran.) This, however, is thought to be a mere Strain upon the text: For the Words are these; That all true Believers shall break their Eggs at the convenient End: and which is the convenient End, seems, in my humble Opinion, to be left to every Man's Conscience, or at least in the Power of the chief Magistrate to determine. Now the Big-Endian Exiles have found so much Credit in the Emperor of Blefuscu's Court; and so much private Assistance and Encouragement from their Party here at home, that a bloody War has been carried on between the two Empires for six and thirty Moons with various Success; during which Time we have lost Forty Capital Ships, and a much greater Number of smaller Vessels, together with thirty thousand of our best Seamen and Soldiers; and the Damage received by the Enemy is reckoned to be somewhat greater than ours. However, they have now equipped a numerous Fleet, and are just preparing to make a Descent upon us: and his Imperial Majesty, placing great Confidence in your Valour and Strength, hath commanded me to lay this Account of his Affairs before you.

### 1.11.3 Effective Address Calculation

An effective address is computed by the processor when executing a Storage Access or Branch instruction (or certain other instructions described in Book II and Book III) when fetching the next sequential instruction, or when invoking a system error handler. The following provides an overview of this process. More detail is provided in the individual instruction descriptions.

Effective address calculations, for both data and instruction accesses, use 64-bit two's complement addition. All 64 bits of each address component participate in the calculation regardless of mode (32-bit or 64-bit). In this computation one operand is an address (which is by definition an unsigned number) and the second is a signed offset. Carries out of the most significant bit are ignored.
In 64-bit mode, the entire 64-bit result comprises the 64 -bit effective address. The effective address arith-
metic wraps around from the maximum address, $2^{64}-1$, to address 0 , except that if the current instruction is at effective address $2^{64}-4$ the effective address of the next sequential instruction is undefined.
In 32-bit mode, the low-order 32 bits of the 64-bit result, preceded by 320 bits, comprise the 64-bit effective address for the purpose of addressing storage. When an effective address is placed into a register by an instruction or event, the value placed into the high-order 32 bits of the register is as follows.

- Load with Update and Store with Update instructions set the high-order 32 bits of register RA to the high-order 32 bits of the 64-bit result.
- In all other cases (e.g., the Link Register when set by Branch instructions having LK=1, Special Purpose Registers when set to an effective address by invocation of a system error handler) the high-order 32 bits of the register
are set to 0s except as described in the last sentence of this paragraph.
As used to address storage, the effective address arithmetic appears to wrap around from the maximum address, $2^{32}-1$, to address 0 , except that if the current instruction is at effective address $2^{32}-4$ the effective address of the next sequential instruction is undefined.
RA is a field in the instruction which specifies an address component in the computation of an effective address. A zero in the RA field indicates the absence of the corresponding address component. A value of zero is substituted for the absent component of the effective address computation. This substitution is shown in the instruction descriptions as (RAIO).
Effective addresses are computed as follows. In the descriptions below, it should be understood that "the contents of a GPR" refers to the entire 64-bit contents, independent of mode, but that in 32-bit mode only bits 32:63 of the 64-bit result of the computation are used to address storage.
- With X-form instructions, in computing the effective address of a data element, the contents of the GPR designated by RB (or the value zero for Iswi and stswi) are added to the contents of the GPR designated by RA or to zero if RA=0 or RA is not used in forming the EA.
- With D-form instructions, the 16-bit $D$ field is sign-extended to form a 64-bit address component. In computing the effective address of a data element, this address component is added to the contents of the GPR designated by RA or to zero if $R A=0$.
- With DS-form instructions, the 14-bit DS field is concatenated on the right with 0b00 and sign-extended to form a 64-bit address component. In computing the effective address of a data element, this address component is added to the contents of the GPR designated by RA or to zero if $R A=0$.

■ With DQ-form instructions, the 12-bit DQ field is concatenated on the right with $0 b 0000$ and sign-extended to form a 64-bit address component. In computing the effective address of a data element, this address component is added to the contents of the GPR designated by RA or to zero if $R A=0$.

■ With I-form Branch instructions, the 24-bit LI field is concatenated on the right with $0 b 00$ and sign-extended to form a 64-bit address component. If $A A=0$, this address component is added to the address of the Branch instruction to form the effective address of the target instruction. If $A A=1$, this address component is the effective address of the target instruction.

- With B-form Branch instructions, the 14-bit BD field is concatenated on the right with ObOO and
sign-extended to form a 64-bit address component. If $A A=0$, this address component is added to the address of the Branch instruction to form the effective address of the target instruction. If $A A=1$, this address component is the effective address of the target instruction.
■ With XL-form Branch instructions, bits 0:61 of the Link Register or the Count Register are concatenated on the right with $0 b 00$ to form the effective address of the target instruction.
- With sequential instruction fetching, the value 4 is added to the address of the current instruction to form the effective address of the next instruction, except that if the current instruction is at the maximum instruction effective address for the mode ( $2^{64}-4$ in 64 -bit mode, $2^{32}-4$ in 32 -bit mode) the effective address of the next sequential instruction is undefined.

If the size of the operand of a Storage Access instruction is more than one byte, the effective address for each byte after the first is computed by adding 1 to the effective address of the preceding byte.

## 28

Power ISA ${ }^{\text {TM }}$ - Book I

## Chapter 2. Branch Facility

### 2.1 Branch Facility Overview

This chapter describes the registers and instructions that make up the Branch Facility.

### 2.2 Instruction Execution Order

In general, instructions appear to execute sequentially, in the order in which they appear in storage. The exceptions to this rule are listed below.

■ Branch instructions for which the branch is taken cause execution to continue at the target address specified by the Branch instruction.
■ Trap instructions for which the trap conditions are satisfied, and System Call and System Call Vectored instructions, cause the appropriate system handler to be invoked.

■ Transaction failure will eventually cause the transaction's failure handler, implied by the tbegin. instruction, to be invoked. See the programming note following the tbegin. description in Section 5.5 of Book II.

- Exceptions can cause the system error handler to be invoked, as described in Section 1.10, "Exceptions" on page 23.

■ Returning from a system service program, system trap handler, or system error handler causes execution to continue at a specified address.

The model of program execution in which the processor appears to execute one instruction at a time, completing each instruction before beginning to execute the next instruction is called the "sequential execution model". In general, the processor obeys the sequential execution model. For the instructions and facilities defined in this Book, the only exceptions to this rule are the following.

- A floating-point exception occurs when the processor is running in one of the Imprecise floating-point exception modes (see Section 4.4). The instruction that causes the exception need not complete before the next instruction begins execution, with
respect to setting exception bits and (if the exception is enabled) invoking the system error handler.
- A Store instruction modifies one or more bytes in an area of storage that contains instructions that will subsequently be executed. Before an instruction in that area of storage is executed, software synchronization is required to ensure that the instructions executed are consistent with the results produced by the Store instruction.


## Programming Note

This software synchronization will generally be provided by system library programs (see Section 1.10 of Book II). Application programs should call the appropriate system library program before attempting to execute modified instructions.

### 2.3 Branch Facility Registers

### 2.3.1 Condition Register

The Condition Register (CR) is a 32-bit register which reflects the result of certain operations, and provides a mechanism for testing (and branching).

| CR |  |  |
| :--- | :---: | :---: |
| 32 |  |  |

## Figure 36. Condition Register

The bits in the Condition Register are grouped into eight 4-bit fields, named CR Field 0 (CRO), ..., CR Field 7 (CR7), which are set in one of the following ways.

■ Specified fields of the CR can be set by a move to the CR from a GPR (mtcrf, mtocrf).

- A specified field of the CR can be set by a move to the CR from another CR field (mcrf), from XER $_{32: 35}$ (mcrxr), or from the FPSCR (mcrfs).
- CR Field 0 can be set as the implicit result of a fixed-point instruction.

■ CR Field 1 can be set as the implicit result of a floating-point instruction.

- CR Field 1 can be set as the implicit result of a decimal floating-point instruction.
- CR Field 6 can be set as the implicit result of a vector instruction.
■ A specified CR field can be set as the result of a Compare instruction or of a tcheck instruction (see Book II).

Instructions are provided to perform logical operations on individual CR bits and to test individual CR bits.

For all fixed-point instructions in which Rc=1, and for addic., andi., and andis., the first three bits of CR Field 0 (bits $32: 34$ of the Condition Register) are set by signed comparison of the result to zero, and the fourth bit of CR Field 0 (bit 35 of the Condition Register) is copied from the SO field of the XER. "Result" here refers to the entire 64-bit value placed into the target register in 64 -bit mode, and to bits $32: 63$ of the 64 -bit value placed into the target register in 32 -bit mode.

```
if (64-bit mode)
    then \(\mathrm{M} \leftarrow 0\)
    else \(\mathrm{M} \leftarrow 32\)
if (target_register) m: \(63<0\) then \(\mathrm{c} \leftarrow 0 \mathrm{~b} 100\)
else if (target_register) \({ }_{\mathrm{M}: 63}>0\) then \(\mathrm{c} \leftarrow 0 \mathrm{~b} 010\)
else \(\quad c \leftarrow 0 \mathrm{~b} 001\)
\(\mathrm{CRO} \leftarrow \mathrm{c} \| \mathrm{XER}_{\mathrm{SO}}\)
```

If any portion of the result is undefined, then the value placed into the first three bits of CR Field 0 is undefined.

The bits of CR Field 0 are interpreted as follows.
Bit Description
$0 \quad$ Negative (LT) The result is negative.
$1 \quad$ Positive (GT)
The result is positive.
2 Zero (EQ)
The result is zero.

## Summary Overflow (SO)

This is a copy of the contents of XER $_{\text {SO }}$ at the completion of the instruction.

With the exception of tcheck, the Transactional Memory instructions set $\mathrm{CRO}_{0: 2}$ indicating the state of the facility prior to instruction execution, or transaction failure. A complete description of the meaning of these bits is given in the instruction descriptions in Section 5.5 of Book II. These bits are interpreted as follows:

| CR0 | Description |
| :--- | :--- |
| 000 II 0 | Transaction state of Non-transactional prior <br> to instruction |
| 010 \|| 0 | Transaction state of Transactional prior to <br> instruction |
| 001 \|| 0 | Transaction state of Suspended prior to <br> instruction <br> $101 ~ \\| ~$ |
| Transaction failure |  |

The tcheck instruction similarly sets bits 1 and 2 of CR field BF to indicate the transaction state, and additionally sets bit 0 to TDOOMED, as defined in Section 5.5 of Book II.

| CR field BF | Description |
| :--- | :--- |
| TDOOMED \|| 00 || 0 | Transaction state of Non-trans- <br> actional prior to instruction <br> TDOOMED \|| 10 \|| 0 |
| Transaction state of Transac- <br> tional prior to instruction |  |
| TDOOMED \|| 01 || 0 | Transaction state of Sus- <br> pended prior to instruction |

> Programming Note
> Setting of bit 3 of the specified CR field to zero by tcheck and of field $\mathrm{CRO}_{3}$ to zero by other TM instructions is intended to preserve these bits for future function. Software should not depend on the bits being zero.

The paste. instruction (see Section 4.4, "Copy-Paste Facility", in Book II) and the stbcx., sthcx., stwcx.,
stdcx., and stqcx. instructions (see Section 4.6.2, "Load and Reserve and Store Conditional Instructions", in Book II) also set CR Field 0.

For all floating-point instructions in which Rc=1, CR Field 1 (bits $36: 39$ of the Condition Register) is set to the Floating-Point exception status, copied from bits 32:35 of the Floating-Point Status and Control Register. This occurs regardless of whether any exceptions are enabled, and regardless of whether the writing of the result is suppressed (see Section 4.4, "Floating-Point Exceptions" on page 132). These bits are interpreted as follows.

## Bit Description

$32 \quad$ Floating-Point Exception Summary (FX)
This is a copy of the contents of FPSCR $_{\text {FX }}$ at the completion of the instruction.
33 Floating-Point Enabled Exception Summary (FEX)
This is a copy of the contents of FPSCR FEX at the completion of the instruction.
$34 \quad$ Floating-Point Invalid Operation Exception Summary (VX)
This is a copy of the contents of FPSCR ${ }_{V x}$ at the completion of the instruction.
Floating-Point Overflow Exception (OX)
This is a copy of the contents of FPSCR OXx at the completion of the instruction.

For Compare instructions, a specified CR field is set to reflect the result of the comparison. The bits of the specified CR field are interpreted as follows. A complete description of how the bits are set is given in the instruction descriptions in Section 3.3.10, "Fixed-Point Compare Instructions" on page 85, and Section 4.6.8, "Floating-Point Compare Instructions" on page 168.

## Bit Description

0 Less Than, Floating-Point Less Than (LT, FL)
For fixed-point Compare instructions, (RA) < SI or (RB) (signed comparison) or (RA) $<{ }^{u}$ UI or (RB) (unsigned comparison). For float-ing-point Compare instructions, (FRA) < (FRB).

1 Greater Than, Floating-Point Greater Than (GT, FG)
For fixed-point Compare instructions, (RA) > SI or (RB) (signed comparison) or (RA) $>{ }^{u} \mathrm{UI}$ or (RB) (unsigned comparison). For float-ing-point Compare instructions, (FRA) > (FRB).

2 Equal, Floating-Point Equal (EQ, FE)
For fixed-point Compare instructions, (RA) = SI, UI, or (RB). For floating-point Compare instructions, $(F R A)=(F R B)$.

3 Summary Overflow, Floating-Point Unordered (SO,FU)
For fixed-point Compare instructions, this is a copy of the contents of XER $_{\text {SO }}$ at the completion of the instruction. For floating-point Compare instructions, one or both of (FRA) and (FRB) is a NaN .

The Vector Integer Compare instructions (see Section 6.9.3, "Vector Integer Compare Instructions") compare two Vector Registers element by element, interpreting the elements as unsigned or signed integers depending on the instruction, and set the corresponding element of the target Vector Register to all 1s if the relation being tested is true and 0 s if the relation being tested is false.

If $R c=1, C R$ Field 6 is set to reflect the result of the comparison, as follows

## Bit Description

$0 \quad$ The relation is true for all element pairs (i.e., VRT is set to all 1s).
$2 \quad$ The relation is false for all element pairs (i.e., VRT is set to all Os).

## 3

0
The Vector Floating-Point Compare instructions compare two Vector Registers word element by word element, interpreting the elements as single-precision floating-point numbers. With the exception of the Vector Compare Bounds Floating-Point instruction, they set the target Vector Register, and CR Field 6 if Rc=1, in the same manner as do the Vector Integer Compare instructions.

## Bit Description

$0 \quad$ The relation is true for all element pairs (i.e., VRT is set to all 1s).

10
2 The relation is false for all element pairs (i.e., VRT is set to all Os).
30
The Vector Compare Bounds Floating-Point instruction on page 331 sets CR Field 6 if $\mathrm{Rc}=1$, to indicate whether the elements in VRA are within the bounds specified by the corresponding element in VRB, as explained in the instruction description. A single-precision floating-point value $x$ is said to be "within the bounds" specified by a single-precision floating-point value $y$ if $-y \leq x \leq y$.

```
Bit Description
O
10
```

Set to indicate whether all four elements in VRA are within the bounds specified by the corresponding element in VRB, otherwise set to 0 .

## 30

### 2.3.2 Link Register

The Link Register (LR) is a 64-bit register. It can be used to provide the branch target address for the Branch Conditional to Link Register instruction, and it holds the return address after Branch instructions for which LK=1 and after System Call Vectored instructions.


Figure 37. Link Register

### 2.3.3 Count Register

The Count Register (CTR) is a 64-bit register. It can be used to hold a loop count that can be decremented during execution of Branch instructions that contain an appropriately coded BO field. If the value in the Count Register is 0 before being decremented, it is -1 afterward. The Count Register can also be used to provide the branch target address for the Branch Conditional to Count Register instruction. The Count Register is modified by the System Call Vectored instruction.


Figure 38. Count Register

### 2.3.4 Target Address Register

The Target Address Register (TAR) is a 64-bit register. It can be used to provide bits 0:61 of the branch target address for the Branch Conditional to Branch Target Address Register instruction. Bits 62:63 are ignored by the hardware but can be set and reset by software.


Figure 39. Target Address Register

## Programming Note

The TAR is reserved for system software.

### 2.4 Branch History Rolling Buffer

The Branch History Rolling Buffer (BHRB) is a buffer containing an implementation-dependent number of entries, referred to as BHRB Entries (BHRBEs), that contain information related to branches that have been taken. Entries are numbered from 0 through n , where n is implementation-dependent but no more than 1023. Entry 0 is the most-recently written entry. The BHRB is read by means of the mfbhrbe instruction.

System software typically controls the availability of the BHRB as well as the number of entries that it contains. If the BHRB is accessed when it is unavailable, the system facility unavailable error handler is invoked.
Various events or actions by the system software may result in the BHRB occasionally being cleared. If BHRB entries are read after this has occurred, Os will be returned. See the description of the mfbhrbe instruction for additional information.

The BHRB is typically used in conjunction with Performance Monitor event-based branches. (See Chapter 7 of Book II.) When used in conjunction with this facility, BESCR $_{\text {PME }}$ is set to 1 to enable Performance Monitor event-based exceptions, and Performance Monitor alerts are enabled to enable the writing of BHRB entries. When a Performance Monitor alert occurs, Performance Monitor alerts are disabled, BHRB entries are no longer written, and an event-based branch occurs. (See Chapter 9 of Book III for additional information on the Performance Monitor.) The event-based branch handler can then access the contents of the BHRB for analysis.

When the BHRB is written by hardware, only those Branch instructions that meet the filtering criteria are written. See Section 9.4.7 of Book III.)

The following paragraphs describe the entries written into the BHRB for various types of Branch instructions for which the branch was taken. In some circumstances, however, the hardware may be unable to make the entry even though the following paragraphs require it. In such cases, the hardware sets the EA field to 0 , and indicates any missed entries using the $T$ and $P$ fields. (See Section 2.4.1.)

When an I-form or B-form Branch instruction is entered into the BHRB, bits 0:61 of the effective address of the Branch instruction are written into the next available entry, except that the entry may or may not be written in the following cases.
■ The effective address of the branch target exceeds the effective address of the Branch instruction by 4.

- The instruction is a B-form Branch, the effective address of the branch target exceeds the effective address of the Branch instruction by 8, and the
instruction immediately following the Branch instruction is not another Branch instruction.

The determination of whether the effective address of the branch target exceeds the effective address of the Branch instruction by 4 or 8 is made modulo $2^{64}$.

## Programming Note

The cases described above, for which the BHRBE need not be written, are cases for which some implementations may optimize the execution of the Branch instruction (first case) or of the Branch instruction and the following instruction (second case) in a manner that makes writing the BHRBE difficult. Such implementations may provide a means by which system software can disable these optimizations, thereby ensuring that the corresponding BHRBEs are written normally.

When an XL-form Branch instruction is entered into the BHRB, bits 0:61 of the effective address of the Branch instruction are written into the next available entry if allowed by the filtering mode; subsequently, bits 0:61 of the effective address of the branch target are written into the following entry.
BHRB entries are written as described above without regard to transactional state and are not removed due to transaction failures.

### 2.4.1 Branch History Rolling Buffer Entry Format

Branch History Rolling Buffer Entries (BHRBEs) have the following format.


Figure 40. Branch History Rolling Buffer Entry
0:61 Effective Address (EA)
When this field is set to a non-zero value, it contains bits 0:61 of the effective address of the instruction indicated by the T field; otherwise this field indicates that the entry is a marker with the meaning specified by the $T$ and P fields.

When the EA field contains a non-zero value, bits 62:63 have the following meanings.

## 62 <br> Target Address ( T )

0 The EA field contains bits 0:61 of the effective address of a Branch instruction for which the branch was taken.
1 The EA field contains bits $0: 61$ of the branch effective address of the branch target of an XL-form Branch instruction for which the branch was taken.
$63 \quad$ Prediction (P)
When $\mathrm{T}=0$, this field has the following meaning.
0 The outcome of the Branch instruction was correctly predicted.
1 The outcome of the Branch instruction was mispredicted.

When $\mathrm{T}=1$, this field has the following meaning.
0 The Branch instruction was predicted to be taken and the target address was predicted correctly, or the target address was not predicted because the branch was predicted to be not taken.
1 The target address was mispredicted.
When the EA field contains a zero value, bits 62:63 specify the type of marker as described below.

## Programming Note

It is expected that programs will not contain Branch instructions with instruction or target effective address equal to 0 . If such instructions exist, programs cannot distinguish between entries that are markers and entries that correspond to instructions with instruction or target effective address 0 .

## Value Meaning

00 This entry either is not implemented or has been cleared. There are no valid entries beyond the current entry.

01-11 Reserved.

### 2.5 Branch Instructions

The sequence of instruction execution can be changed by the Branch instructions. Because all instructions are on word boundaries, bits 62 and 63 of the generated branch target address are ignored by the processor in performing the branch.

The Branch instructions compute the effective address (EA) of the target in one of the following five ways, as described in Section 1.11.3, "Effective Address Calculation" on page 27.

1. Adding a displacement to the address of the Branch instruction (Branch or Branch Conditional with $A A=0$ ).
2. Specifying an absolute address (Branch or Branch Conditional with $\mathrm{AA}=1$ ).
3. Using the address contained in the Link Register (Branch Conditional to Link Register).
4. Using the address contained in the Count Register (Branch Conditional to Count Register).
5. Using the address contained in the Target Address Register (Branch Conditional to Target Address Register).
In all five cases, in 32-bit mode the final step in the address computation is setting the high-order 32 bits of the target address to 0 .

For the first two methods, the target addresses can be computed sufficiently ahead of the Branch instruction that instructions can be prefetched along the target path. For the third through fifth methods, prefetching instructions along the target path is also possible provided the Link Register or the Count Register is loaded sufficiently ahead of the Branch instruction.

Branching can be conditional or unconditional, and the return address can optionally be provided. If the return address is to be provided ( $\mathrm{LK}=1$ ), the effective address of the instruction following the Branch instruction is placed into the Link Register after the branch target address has been computed; this is done regardless of whether the branch is taken.

For Branch Conditional instructions, the BO field specifies the conditions under which the branch is taken, as shown in Figure 41. In the figure, $M=0$ in 64 -bit mode and $M=32$ in 32 -bit mode.

| BO | Description |
| :---: | :---: |
| 0000z | Decrement the CTR, then branch if the decremented $\mathrm{CTR}_{\mathrm{M}: 63} \neq 0$ and $\mathrm{CR}_{\mathrm{BI}}=0$ |
| 0001z | Decrement the CTR, then branch if the decremented $\mathrm{CTR}_{\mathrm{M}: 63}=0$ and $\mathrm{CR}_{\mathrm{BI}}=0$ |
| 001at | Branch if $\mathrm{CR}_{\mathrm{BI}}=0$ |
| 0100z | Decrement the CTR, then branch if the decremented $\mathrm{CTR}_{\mathrm{M}: 63} \neq 0$ and $\mathrm{CR}_{\mathrm{BI}}=1$ |
| 0101z | Decrement the CTR, then branch if the decremented $\mathrm{CTR}_{\mathrm{M}: 63}=0$ and $\mathrm{CR}_{\mathrm{BI}}=1$ |
| 011at | Branch if $\mathrm{CR}_{\mathrm{Bl}}=1$ |
| 1a00t | Decrement the CTR, then branch if the decremented CTR $_{\mathrm{M}: 63} \neq 0$ |
| 1a01t | Decrement the CTR, then branch if the decremented CTR $_{\text {M: } 63}=0$ |
| 1z1zz | Branch always |
| Notes: <br> 1. "z" denotes a bit that is ignored. <br> 2. The "a" and " $t$ " bits are used as described below. |  |

Figure 41. BO field encodings
The " $a$ " and " t " bits of the BO field can be used by software to provide a hint about whether the branch is likely to be taken or is likely not to be taken, as shown in Figure 42.

| at | Hint |
| :--- | :--- |
| 00 | No hint is given |
| 01 | Reserved |
| 10 | The branch is very likely not to be taken |
| 11 | The branch is very likely to be taken |

Figure 42. "at" bit encodings

## Programming Note

Many implementations have dynamic mechanisms for predicting whether a branch will be taken. Because the dynamic prediction is likely to be very accurate, and is likely to be overridden by any hint provided by the "at" bits, the "at" bits should be set to Ob00 unless the static prediction implied by at=0b10 or at=0b11 is highly likely to be correct.

For Branch Conditional to Link Register, Branch Conditional to Count Register, and Branch Conditional to Target Address Register instructions, the BH field provides
a hint about the use of the instruction, as shown in Figure 43.

| BH | Hint |
| :---: | :---: |
| 00 | bclr[ [] : The instruction is a subroutine return <br> $\boldsymbol{b c c t r}[\Pi]$ and $\boldsymbol{b c t a r}[]$ The instruction is not a subroutine return; the target address is likely to be the same as the target address used the preceding time the branch was taken |
| 01 | bclr[ $[$ : The instruction is not a subroutine return; the target address is likely to be the same as the target address used the preceding time the branch was taken <br> $\boldsymbol{b c c t r}[1]$ and $\boldsymbol{b c t a r}[I]:$ Reserved |
| 10 | Reserved |
| 11 | bclı[I], bcctr[[], and bctar[l]: The target address is not predictable |

Figure 43. BH field encodings

## Programming Note

The hint provided by the BH field is independent of the hint provided by the "at" bits (e.g., the BH field provides no indication of whether the branch is likely to be taken).

## Extended mnemonics for branches

Many extended mnemonics are provided so that Branch Conditional instructions can be coded with portions of the BO and BI fields as part of the mnemonic rather than as part of a numeric operand. Some of these are shown as examples with the Branch instructions. See Appendix $C$ for additional extended mnemonics.

## Programming Note

The hints provided by the "at" bits and by the BH field do not affect the results of executing the instruction.

The "z" bits should be set to 0 , because they may be assigned a meaning in some future version of the architecture.

## Programming Note

Many implementations have dynamic mechanisms for predicting the target addresses of bclr $[/]$ and $\operatorname{bcctr}[I]$ instructions. These mechanisms may cache return addresses (i.e., Link Register values set by Branch instructions for which LK=1 and for which the branch was taken, other than the special form shown in the first example below) and recently used branch target addresses. To obtain the best performance across the widest range of implementations, the programmer should obey the following rules.

- Use Branch instructions for which LK=1 only as subroutine calls (including function calls, etc.), or in the special form shown in the first example below.
- Pair each subroutine call (i.e., each Branch instruction for which $L K=1$ and the branch is taken, other than the special form shown in the first example below) with a bclr instruction that returns from the subroutine and has $\mathrm{BH}=0 \mathrm{~b} 00$.
- Do not use bcIrl as a subroutine call. (Some implementations access the return address cache at most once per instruction; such implementations are likely to treat bcIrl as a subroutine return, and not as a subroutine call.)
- For bclr[ [] and bcctr[ [] , use the appropriate value in the BH field.
The following are examples of programming conventions that obey these rules. In the examples, BH is assumed to contain 0b00 unless otherwise stated. In addition, the "at" bits are assumed to be coded appropriately.
Let $A, B$, and Glue be specific programs.
- Obtaining the address of the next instruction: Use the following form of Branch and Link. bcl 20,31,\$+4
- Loop counts:

Keep them in the Count Register, and use a bc instruction (LK=0) to decrement the count and to branch back to the beginning of the loop if the decremented count is nonzero.

- Computed goto's, case statements, etc.:

Use the Count Register to hold the address to
branch to, and use a bcctr instruction (LK=0, and $\mathrm{BH}=0 \mathrm{~b} 11$ if appropriate) to branch to the selected address.

- Direct subroutine linkage:

Here $A$ calls $B$ and $B$ returns to $A$. The two branches should be as follows.

- A calls B: use a blor bclinstruction (LK=1).
- B returns to A: use a bclr instruction (LK=0) (the return address is in, or can be restored to, the Link Register).
- Indirect subroutine linkage:

Here A calls Glue, Glue calls B, and B returns to A rather than to Glue. (Such a calling sequence is common in linkage code used when the subroutine that the programmer wants to call, here $B$, is in a different module from the caller; the Binder inserts "glue" code to mediate the branch.) The three branches should be as follows.

- A calls Glue: use a bl or bcl instruction ( $\mathrm{LK}=1$ ).
- Glue calls B: place the address of B into the Count Register, and use a bcctr instruction (LK=0).
- B returns to A: use a bclr instruction (LK=0) (the return address is in, or can be restored to, the Link Register).
- Function call:

Here A calls a function, the identity of which may vary from one instance of the call to another, instead of calling a specific program B. This case should be handled using the conventions of the preceding two bullets, depending on whether the call is direct or indirect, with the following differences.

- If the call is direct, place the address of the function into the Count Register, and use a $\boldsymbol{b} \boldsymbol{c} \boldsymbol{c t r l}$ instruction (LK=1) instead of a blor bcl instruction.
- For the bcctr[I] instruction that branches to the function, use $\mathrm{BH}=0 \mathrm{~b} 11$ if appropriate.


## Compatibility Note

The bits corresponding to the current "a" and " t " bits, and to the current " $z$ " bits except in the "branch always" BO encoding, had different meanings in versions of the architecture that precede Version 2.00.

■ The bit corresponding to the " t " bit was called the " $y$ " bit. The " $y$ " bit indicated whether to use the architected default prediction ( $y=0$ ) or to use the complement of the default prediction ( $y=1$ ). The default prediction was defined as follows.

- If the instruction is $\boldsymbol{b c}[/ \Omega[\boldsymbol{a}]$ with a negative value in the displacement field, the branch is taken. (This is the only case in which the prediction corresponding to the " y " bit differs from the prediction corresponding to the " t " bit.)
- In all other cases (bc[ $\boldsymbol{I}[a]$ with a nonnegative value in the displacement field, bclr[]], or $\boldsymbol{b c c t r}[\mathrm{I}]$ ), the branch is not taken.
- The BO encodings that test both the Count Register and the Condition Register had a " y " bit in place of the current " $z$ " bit. The meaning of the " $y$ " bit was as described in the preceding item.
■ The "a" bit was a "z" bit.
Because these bits have always been defined either to be ignored or to be treated as hints, a given program will produce the same result on any implementation regardless of the values of the bits. Also, because even the " $y$ " bit is ignored, in practice, by most processors that comply with versions of the architecture that precede Version 2.00, the performance of a given program on those processors will not be affected by the values of the bits.


## Branch

| b | target_addr |
| :--- | :--- |
| ba | target_addr |
| bl | target_addr |
| bla | target_addr |

target_addr specifies the branch target address.
If $A A=0$ then the branch target address is the sum of $\mathrm{LI} \| \mathrm{ObOO}$ sign-extended and the address of this instruction, with the high-order 32 bits of the branch target address set to 0 in 32-bit mode.
If $A A=1$ then the branch target address is the value $\mathrm{LI} \| \mathrm{Ob00}$ sign-extended, with the high-order 32 bits of the branch target address set to 0 in 32-bit mode.
If $L K=1$ then the effective address of the instruction fol-
If $\mathrm{LK}=1$ then the effective address of the instruction fol-
lowing the Branch instruction is placed into the Link Register.

## Special Registers Altered:

LR

```
if AA then NIA \(\leftarrow_{\text {iea }} \operatorname{EXTS}(L I| | 0 b 00)\)
else \(\quad\) NIA \(\leftarrow_{\text {iea }}\) CIA \(+\operatorname{EXTS}(L I \| 0 b 00)\)
if \(L K\) then \(L R \leftarrow{ }_{\text {iea }}\) CIA +4
NIA }\mp@subsup{\leftarrow}{iea}{\mathrm{ CIA + EXTS(LI | Ob00)}
```

(if $\mathrm{LK}=1$ )

I-form
( $\mathrm{AA}=0 \mathrm{LK}=0$ )
( $\mathrm{A} A=1 \mathrm{LK}=0$ )
( $\mathrm{A} A=0 \mathrm{LK}=1$ )
( $\mathrm{A} A=1 \mathrm{LK}=1$ )

| 18 | LI | AA | LK |  |
| :---: | :---: | :---: | :---: | :---: |
| 0 | 6 |  | 30 | 31 |

## Branch Conditional

| bc <br> bca <br> bcl <br> bcla | BO,BI,target_addr |  |  |  | ( $\mathrm{AA}=0 \mathrm{LK}=0$ ) |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | BO,BI,target_addr |  |  |  | ( $\mathrm{AA}=1 \mathrm{LK}=0$ ) |  |
|  | BO,BI,target_addr |  |  |  | (AA=0 LK=1) |  |
|  | BO,BI,target_addr |  |  |  | 1 LK | K=1) |
| 16 | BO | BI |  | BD |  | LK |
| 0 | 6 | 11 | 16 |  |  |  |

$$
\begin{aligned}
& \text { if (64-bit mode) } \\
& \text { then } \mathrm{M} \leftarrow 0 \\
& \text { else } \mathrm{M} \leftarrow 32 \\
& \text { if } \neg \mathrm{BO}_{2} \text { then } \mathrm{CTR} \leftarrow \mathrm{CTR}-1 \\
& \text { ctr_ok } \leftarrow \mathrm{BO}_{2} \mid\left(\left(\mathrm{CTR}_{\mathrm{M}: 63} \neq 0\right) \oplus \mathrm{BO}_{3}\right) \\
& \text { cond_ok } \leftarrow \mathrm{BO}_{0} \mid\left(\mathrm{CR}_{\mathrm{BI}+32} \equiv \mathrm{BO}_{1}\right) \\
& \text { if ctr_ok \& cond_ok then } \\
& \begin{array}{l}
\text { if AA then NIA } \leftarrow_{\text {iea }} \operatorname{EXTS}(\mathrm{BD}| | 0 \mathrm{DOO}) \\
\text { else NIA } \leftarrow_{\text {iea }} \text { CIA }+\operatorname{EXTS}(\mathrm{BD} \| 0 \mathrm{l} 00) \\
\text { if LK then } \mathrm{LR} \leftarrow_{\text {iea }} \text { CIA }+4
\end{array}
\end{aligned}
$$

$\mathrm{BI}+32$ specifies the Condition Register bit to be tested. The BO field is used to resolve the branch as described in Figure 41. target_addr specifies the branch target address.

If $A A=0$ then the branch target address is the sum of $B D \| 0 b 00$ sign-extended and the address of this instruction, with the high-order 32 bits of the branch target address set to 0 in 32-bit mode.
If $A A=1$ then the branch target address is the value BD || 0b00 sign-extended, with the high-order 32 bits of the branch target address set to 0 in 32-bit mode.

If $\mathrm{LK}=1$ then the effective address of the instruction following the Branch instruction is placed into the Link Register.
Special Registers Altered:
CTR
(if $\mathrm{BO}_{2}=0$ )
LR
(if $\mathrm{LK}=1$ )

## Extended Mnemonics:

Examples of extended mnemonics for Branch Conditional:

| Extended: |  | Equivalent to: |  |
| :--- | :--- | :--- | ---: |
| blt | target | bc | 12,0, target |
| bne | cr2,target | bc | 4,10, target |
| bdnz | target | bc | 16,0, target |

## Branch Conditional to Link Register



if (64-bit mode)
then $\mathrm{M} \leftarrow 0$
else $M \leftarrow 32$
if $\neg \mathrm{BO}_{2}$ then $\mathrm{CTR} \leftarrow \mathrm{CTR}-1$
ctr_ok $\leftarrow \mathrm{BO}_{2} \mid\left(\left(\mathrm{CTR}_{\mathrm{M}: 63} \neq 0\right) \oplus \mathrm{BO}_{3}\right.$
cond_ok $\leftarrow \mathrm{BO}_{0} \mid \quad\left(\mathrm{CR}_{\mathrm{BI}+32} \equiv \mathrm{BO}_{1}\right)$
if ctr_ok \& cond_ok then NIA $\leftarrow_{\text {iea }} \mathrm{LR}_{0: 61}| | 0 \mathrm{~b} 00$
if LK then LR $\leftarrow_{\text {iea }}$ CIA +4
$\mathrm{BI}+32$ specifies the Condition Register bit to be tested. The BO field is used to resolve the branch as described in Figure 41. The BH field is used as described in Figure 43. The branch target address is $\mathrm{LR}_{0: 61}$ II $0 b 00$, with the high-order 32 bits of the branch target address set to 0 in 32 -bit mode.

If $\mathrm{LK}=1$ then the effective address of the instruction following the Branch instruction is placed into the Link Register.

## Special Registers Altered:

| CTR | (if $\mathrm{BO}_{2}=0$ ) |
| :--- | ---: |
| LR | (if $\mathrm{LK}=1$ ) |

## Extended Mnemonics:

Examples of extended mnemonics for Branch Conditional to Link Register:

| Extended: |  | Equivalent to: |  |
| :--- | :--- | :--- | :--- |
| bclr | 4,6 | bclr | $4,6,0$ |
| bltlr |  | bclr | $12,0,0$ |
| bnelr | cr2 | bclr | $4,10,0$ |
| bdnzlr |  | bclr | $16,0,0$ |

## Programming Note

bclr, bclrl, bcctr, and bcctrl each serve as both a basic and an extended mnemonic. The Assembler will recognize a bclr, bclrl, bcctr, or bcctrl mnemonic with three operands as the basic form, and a bclr, bclrl, bcctr, or bcctrl mnemonic with two operands as the extended form. In the extended form the BH operand is omitted and assumed to be Ob00.
Branch Conditional to Count Register
XL-form

| $\begin{array}{ll} 19 \\ 0 \end{array}$ | $6 \mathrm{BO}$ | ${ }_{11} \mathrm{BI}$ | $1{ }_{16} / 1 /$ | $\begin{aligned} & \mathrm{BH} \\ & 19 \end{aligned}$ | 21 | 528 | LK 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

cond_ok $\leftarrow \mathrm{BO}_{0} \mid \quad\left(\mathrm{CR}_{\mathrm{BI}+32} \equiv \mathrm{BO}_{1}\right)$
if cond_ok then NIA $\leftarrow_{\text {iea }} \mathrm{CTR}_{0: 61}| | 0 \mathrm{bb00}$
if LK then LR $\leftarrow_{\text {iea }}$ CIA +4
$\mathrm{BI}+32$ specifies the Condition Register bit to be tested. The BO field is used to resolve the branch as described in Figure 41. The BH field is used as described in Figure 43. The branch target address is $\mathrm{CTR}_{0: 61}$ II Ob00, with the high-order 32 bits of the branch target address set to 0 in 32-bit mode.
If $\mathrm{LK}=1$ then the effective address of the instruction following the Branch instruction is placed into the Link Register.
If the "decrement and test CTR" option is specified $\left(\mathrm{BO}_{2}=0\right)$, the instruction form is invalid.

## Special Registers Altered:

LR
(if $\mathrm{LK}=1$ )

## Extended Mnemonics:

Examples of extended mnemonics for Branch Conditional to Count Register.

| Extended: |  |
| :--- | :--- |
| bcctr | 4,6 |
| bltctr |  |
| bnectr | cr2 |

Equivalent to:
bcctr 4,6,0
bcctr 12,0,0
bectr 4,10,0

## Version 3.0

## Branch Conditional to Branch Target Address Register

|  | XL-form |  |
| :--- | ---: | ---: |
|  |  | $($ LLK $=0)$ |
| bctar | BO,BI,BH | $(L K=1)$ |


| $19$ | $6$ |  | 16 | $\left\lvert\, \begin{aligned} & \mathrm{BH} \\ & 19\end{aligned}\right.$ | 21 | 560 | LK 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

```
if (64-bit mode)
    then M}\leftarrow
    else M}\leftarrow3
if }\neg\mp@subsup{\textrm{BO}}{2}{}\mathrm{ then CTR }\leftarrow\textrm{CTR - 1
ctr_ok \leftarrow BO2 | ((CTR M:63 = 0) \oplus BO
```



```
if ctr_ok & cond_ok then NIA \leftarrow &iea TAR O:61 || 0b00
if LK then LR }\mp@subsup{\leftarrow}{iea}{CIA + 4
```

$\mathrm{Bl}+32$ specifies the Condition Register bit to be tested. The BO field is used to resolve the branch as described in Figure 41. The BH field is used as described in Figure 43. The branch target address is $\mathrm{TAR}_{0: 61}$ II ObOO, with the high-order 32 bits of the branch target address set to 0 in 32-bit mode.

If $\mathrm{LK}=1$ then the effective address of the instruction following the Branch instruction is placed into the Link Register.

## Special Registers Altered:

| CTR | (if $\mathrm{BO}_{2}=0$ ) |
| :--- | ---: |
| LR | (if $\mathrm{LK}=1$ ) |

## Programming Note

In some systems, the system software will restrict usage of the bctar[l] instruction to only selected programs. If an attempt is made to execute the instruction when it is not available, the system error handler will be invoked. See Book III for additional information.

### 2.6 Condition Register Instructions

### 2.6.1 Condition Register Logical Instructions

The Condition Register Logical instructions have preferred forms; see Section 1.9.1. In the preferred forms, the BT and BB fields satisfy the following rule.

- The bit specified by BT is in the same Condition Register field as the bit specified by BB.


## Extended mnemonics for Condition Register logical operations

A set of extended mnemonics is provided that allow additional Condition Register logical operations, beyond those provided by the basic Condition Register Logical instructions, to be coded easily. Some of these are shown as examples with the Condition Register Logical instructions. See Appendix C for additional extended mnemonics.


The bit in the Condition Register specified by BA+32 is ANDed with the bit in the Condition Register specified by $\mathrm{BB}+32$, and the result is placed into the bit in the Condition Register specified by BT+32.

Special Registers Altered:
$\mathrm{CR}_{\mathrm{BT}+32}$

Condition Register OR XL-form
cror $\quad \mathrm{BT}, \mathrm{BA}, \mathrm{BB}$

| 19 | BT |  | BA | BB |  | 449 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

$\mathrm{CR}_{\mathrm{BT}+32} \leftarrow \mathrm{CR}_{\mathrm{BA}+32} \mid \mathrm{CR}_{\mathrm{BB}+32}$
The bit in the Condition Register specified by BA+32 is ORed with the bit in the Condition Register specified by $\mathrm{BB}+32$, and the result is placed into the bit in the Condition Register specified by BT+32.
Special Registers Altered:
$\mathrm{CR}_{\mathrm{BT}+32}$
Extended Mnemonics:
Example of extended mnemonics for Condition Register OR:

| Extended: | Equivalent to: |
| :--- | :--- |
| crmove $\mathrm{Bx}, \mathrm{By}$ | cror $\quad \mathrm{Bx}, \mathrm{By}, \mathrm{By}$ |

The bit in the Condition Register specified by BA+32 is ANDed with the bit in the Condition Register specified by $\mathrm{BB}+32$, and the complemented result is placed into the bit in the Condition Register specified by BT+32.

## Special Registers Altered:

$\mathrm{CR}_{\mathrm{BT}+32}$

Condition Register XOR XL-form
crxor BT,BA,BB

| 19 | BT | BA | BB | 193 | / |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 | 31 |

$\mathrm{CR}_{\mathrm{BT}+32} \leftarrow \mathrm{CR}_{\mathrm{BA}+32} \oplus \mathrm{CR}_{\mathrm{BB}+32}$
The bit in the Condition Register specified by BA+32 is XORed with the bit in the Condition Register specified by $\mathrm{BB}+32$, and the result is placed into the bit in the Condition Register specified by BT+32.
Special Registers Altered:
$\mathrm{CR}_{\mathrm{BT}+32}$

## Extended Mnemonics:

Example of extended mnemonics for Condition Register XOR:

$$
\begin{aligned}
& \text { Extended: Equivalent to: } \\
& \text { crclr } \mathrm{Bx} \quad \mathrm{crxor} \mathrm{Bx}, \mathrm{Bx}, \mathrm{Bx}
\end{aligned}
$$

| Condition Register NOR |
| :--- |
| crnor |
| BT,BA,BB |
| 19  BT  XL-form |

$\mathrm{CR}_{\mathrm{BT}+32} \leftarrow \neg\left(\mathrm{CR}_{\mathrm{BA}+32} \mid \mathrm{CR}_{\mathrm{BB}+32}\right)$
The bit in the Condition Register specified by BA+32 is ORed with the bit in the Condition Register specified by $\mathrm{BB}+32$, and the complemented result is placed into the bit in the Condition Register specified by BT+32.

## Special Registers Altered:

$$
\mathrm{CR}_{\mathrm{BT}+32}
$$

## Extended Mnemonics:

Example of extended mnemonics for Condition Register NOR:

## Extended: Equivalent to: crnot Bx,By crnor Bx,By,By <br> Condition Register AND with Complement XL-form

$$
\text { crandc } \quad B T, B A, B B
$$

| 19 | BT | BA | BB |  | 129 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |

$$
\mathrm{CR}_{\mathrm{BT}+32} \leftarrow \mathrm{CR}_{\mathrm{BA}+32} \& \neg \mathrm{CR}_{\mathrm{BB}+32}
$$

The bit in the Condition Register specified by BA+32 is ANDed with the complement of the bit in the Condition Register specified by $\mathrm{BB}+32$, and the result is placed into the bit in the Condition Register specified by $B T+32$.
Special Registers Altered:
$\mathrm{CR}_{\mathrm{BT}+32}$

Condition Register Equivalent
XL-form
creqv $\quad B T, B A, B B$

| 19 | BT | BA | BB |  | 289 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

$\mathrm{CR}_{\mathrm{BT}+32} \leftarrow \mathrm{CR}_{\mathrm{BA}+32} \equiv \mathrm{CR}_{\mathrm{BB}+32}$
The bit in the Condition Register specified by BA+32 is XORed with the bit in the Condition Register specified by $\mathrm{BB}+32$, and the complemented result is placed into the bit in the Condition Register specified by BT+32.

## Special Registers Altered:

$$
\mathrm{CR}_{\mathrm{BT}+32}
$$

Extended Mnemonics:
Example of extended mnemonics for Condition Register Equivalent.

$$
\begin{array}{ll}
\text { Extended: } & \text { Equivalent to: } \\
\text { crset } \mathrm{Bx} & \text { creqv } \mathrm{Bx}, \mathrm{Bx}, \mathrm{Bx}
\end{array}
$$

## Condition Register OR with Complement XL-form

crorc $\quad B T, B A, B B$

| 19 |  | BT | BA | BB |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |

$$
\mathrm{CR}_{\mathrm{BT}+32} \leftarrow \mathrm{CR}_{\mathrm{BA}+32} \mid \neg \mathrm{CR}_{\mathrm{BB}+32}
$$

The bit in the Condition Register specified by BA+32 is ORed with the complement of the bit in the Condition Register specified by $\mathrm{BB}+32$, and the result is placed into the bit in the Condition Register specified by BT+32.
Special Registers Altered:
$\mathrm{CR}_{\mathrm{BT}+32}$

### 2.6.2 Condition Register Field Instruction

## Move Condition Register Field

XL-form
morf BF,BFA

| 19 | BF | // | BFA | I/ | I/I |  | 0 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  | 11 | 14 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |  |  |

$\mathrm{CR}_{4 \times \mathrm{BF}+32: 4 \times \mathrm{BF}+35} \leftarrow \mathrm{CR}_{4 \times \mathrm{BFA}+32: 4 \times \mathrm{BFA}+35}$
The contents of Condition Register field BFA are copied to Condition Register field BF.
Special Registers Altered:
CR field BF

### 2.7 System Call Instructions

These instructions provide the means by which a program can call upon the system to perform a service.

## System Call SC-form

## sc LEV

| 17 | $/ / /$ | $/ / /$ | LEV | LEV | // | 1 | $/$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  | 11 | 16 | 20 |  |

## System Call Vectored

SC-form
scv LEV

| 17 |  | I/I |  | I/I | I/ | LEV | I/ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  | 11 | 16 | 20 | 27 |
| 30 | 31 |  |  |  |  |  |  |

These instructions call the system to perform a service. A complete description of these instructions can be found in Section 3.3.1 of Book III.

The first form of the instruction (sc) provides a single system call. The second form of the instruction (scv) provides the capability for 128 unique system calls.

The use of the LEV field is described in Book III. In the first form of the instruction the LEV values greater than 1 are reserved, and bits 0:5 of the LEV field (instruction bits 20:25) are treated as a reserved field.

When control is returned to the program that executed
| the System Call or System Call Vectored instruction, the contents of the registers will depend on the register conventions used by the program providing the system service.

These instructions are context synchronizing (see Book III).

I

## Special Registers Altered:

Dependent on the system service

## Programming Note

$\boldsymbol{s c}$ serves as both a basic and an extended mnemonic. The Assembler will recognize an sc mnemonic with one operand as the basic form, and an sc mnemonic with no operand as the extended form. In the extended form the LEV operand is omitted and assumed to be 0 .

In application programs the value of the LEV operand for $\boldsymbol{s c}$ should be 0.

### 2.8 Branch History Rolling Buffer Instructions

The Branch History Rolling Buffer instructions enable application programs to clear and read the BHRB. The availability of these instructions is controlled by the system software. (See Chapter 9 of Book III.) When an attempt is made to execute these instructions when

## Clear BHRB

X-form
clrbhrb

| 31 | $/ / /$ |  | $/ / /$ | $/ / /$ |  | 430 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 | 16 | 21 |  |

for $\mathrm{n}=0$ to (number_of_BHRBEs implemented - 1) $\operatorname{BHRB}(\mathrm{n}) \leftarrow 0$

All BHRB entries are set to Os.
Special Registers Altered:
None.
they are unavailable, the system facility unavailable error handler is invoked.

Move From Branch History Rolling Buffer Entry

XFX-form
mfbhrbe RT,BHRBE

| 31 | RT | BHRBE |  |  | 302 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 |  |  | 11 |  |

$$
\mathrm{n} \leftarrow \text { BHRBE }_{0: 9}
$$

If $n<$ number of BHRBES implemented then
$\mathrm{RT} \leftarrow \operatorname{BHRBE}(\mathrm{n})$
else
$\mathrm{RT} \leftarrow{ }^{64} 0$
The BHRBE field denotes an entry in the BHRB. If the designated entry is within the range of BHRB entries implemented and Performance Monitor alterts are disable (see Section 9.5 of Book III), the contents of the designated BHRB entry are placed into register RT; otherwise, ${ }^{64} 0$ s are placed into register RT.
In order to ensure that the current BHRB contents are read by this instruction, one of the following must have occurred prior to this instruction and after all previous Branch and clrbhrb instructions have completed.

- an event-based branch has occurred

■ an rfebb (see Chapter 7 of Book II) has been executed

- a context synchronizing event (see Section 1.5 of Book III) other than isynch Section 4.6.1 of Book II) has occurred.


## Special Registers Altered:

None

## Programming Note

In order to read all the BHRB entries containing information about taken branches, software should read the entries starting from entry number 0 and continuing until an entry containing all 0 s is read or until all implemented BHRB entries have been read.

Since the number of BHRB entries may decrease or the BHRB may be cleared at any time, if a given entry, $m$, is read as not containing all 0 s and is read again subsequently, the subsequent read may return all Os even though the program has not executed clrbhrb.

## Chapter 3. Fixed-Point Facility

### 3.1 Fixed-Point Facility Overview

This chapter describes the registers and instructions that make up the Fixed-Point Facility.

### 3.2 Fixed-Point Facility Registers

### 3.2.1 General Purpose Registers

All manipulation of information is done in registers internal to the Fixed-Point Facility. The principal storage internal to the Fixed-Point Facility is a set of 32 General Purpose Registers (GPRs). See Figure 44.

| GPR 0 |
| :---: |
| GPR 1 |
| $\ldots$ |
| $\ldots$ |
| GPR 30 |
| GPR 31 |
| 0 |

Figure 44. General Purpose Registers
Each GPR is a 64-bit register.

### 3.2.2 Fixed-Point Exception Register

The Fixed-Point Exception Register (XER) is a 64-bit register.


Figure 45. Fixed-Point Exception Register
The bit definitions for the Fixed-Point Exception Register are shown below. Here $M=0$ in 64 -bit mode and $\mathrm{M}=32$ in 32-bit mode.

The bits are set based on the operation of an instruction considered as a whole, not on intermediate results (e.g., the Subtract From Carrying instruction, the result of which is specified as the sum of three values, sets bits in the Fixed-Point Exception Register based on the entire operation, not on an intermediate sum).

## Bit(s Description

0:31 Reserved
Summary Overflow (SO)
The Summary Overflow bit is set to 1 whenever an instruction (except mtspr) sets the Overflow bit. Once set, the SO bit remains set until it is cleared by an mtspr instruction (specifying the XER) or an merxr instruction. It is not altered by Compare instructions, or by other instructions (except mtspr to the XER, and mcrxr) that cannot overflow. Executing an mtspr instruction to the XER, supplying the values 0 for SO and 1 for OV, causes SO to be set to 0 and $O V$ to be set to 1 .

## Overflow (OV)

The Overflow bit is set to indicate that an overflow has occurred during execution of an instruction.
XO-form Add, Subtract From, and Negate instructions having $\mathrm{OE}=1$ set it to 1 if the carry out of bit M is not equal to the carry out of bit $\mathrm{M}+1$, and set it to 0 otherwise.
XO-form Multiply Low and Divide instructions having $\mathrm{OE}=1$ set it to 1 if the result cannot be represented in 64 bits (mulld, divd, divde, divdu, divdeu) or in 32 bits (mullw, divw, divwe, divwu, divweu), and set it to 0 otherwise. The OV bit is not altered by Compare
instructions, or by other instructions (except $\boldsymbol{m t s p r}$ to the XER, and merxr) that cannot overflow.

Carry (CA)
The Carry bit is set as follows, during execution of certain instructions. Add Carrying, Subtract From Carrying, Add Extended, and Subtract From Extended types of instructions set it to 1 if there is a carry out of bit M , and set it to 0 otherwise. Shift Right Algebraic instructions set it to 1 if any 1 -bits have been shifted out of a negative operand, and set it to 0 otherwise. The CA bit is not altered by Compare instructions, or by other instructions (except Shift Right Algebraic, mtspr to the XER, and merxr) that cannot carry.
35:43 Reserved
44 Overflow32 (OV32)
OV32 is set whenever OV is set, and is set to the same value that OV is defined to be set to in 32-bit mode.
$45 \quad$ Carry 32 (CA32)
CA32 is set whenever CA is set, and is set to the same value that CA is defined to be set to in 32-bit mode.

46:56 Reserved
57:63 This field specifies the number of bytes to be transferred by a Load String Indexed or Store String Indexed instruction.

### 3.2.3 VR Save Register

| VRSAVE |  |
| :--- | :--- |
| 32 | 63 |

The VR Save Register (VRSAVE) is a 32-bit register that can be used as a software use SPR; see Section 6.3.3.

### 3.2.4 Load Monitored Region Registers

The Load Monitored Region registers are used to specify a Load Monitored region, and to specify the sections of the region that are enabled.

The Load Monitored region is a contiguous region of storage specified by the Load Monitored Region register. Load Monitored regions range in size from 32 MB to 1 TB. All regions are powers of 2 in length and are aligned (see Section 1.11.1).
The Load Monitored region is divided into 64 sections of equal length, each of which can be separately enabled. The Load Monitored Section Enable Register specifies which sections are enabled.

When the value read by an Idmx instruction is equal to an effective address within an enabled section of the Load Monitored region specified by these registers, a Load Monitored event-based exception occurs if BESCR $_{\text {GE LME }}=(1,1)$. (See Section 7.2.1 of Book II.)

### 3.2.4.1 Load Monitored Region Register

The Load Monitored Region Register (LMRR) specifies the base address and size of the Load Monitored region.

| Base Effective Address | $/ / /$ | Size |
| :--- | :--- | :--- |
| 0 | 39 | 60 |

Figure 46. Load Monitored Region Register

## 0:38 Base Effective Address

This field specifies the most-significant bits of the effective address of the first byte in the Load Monitored region. Only the high-order $39-S$ bits of the field are used, where $S$ is the value specified in the Size field; the remaining S+35 bits of the effective address are assumed to be Os. For example, for a 32 MB region all 39 bits of the field are used, and for a 1 TB region only bits 0:23 are used.

## 39:59 Reserved

60:63 Size
This field specifies the size of the Load Monitored region. The size of the Load Monitored region is $2^{(\mathrm{S}+25)}$ bytes, where $S$ is the value in the Size field. For example, $S=0$ specifies a 32 MB Load Monitored region, $S=1$ specifies a 64 $M B$ region, and $S=15$ specifies a 1 TB region.

### 3.2.4.2 Load Monitored Section Enable Register

The Load Monitored Section Enable Register (LMSER) specifies the sections of the Load Monitored region that are enabled. The sections are numbered $0-63$, where bit n corresponds to section number n .


Figure 47. Load Monitored Section Enable Register
Each bit, n, of the Load Monitored Section Enable Register specifies whether or not section number $n$ is enabled. If bit n is 0 , section n is disabled; if bit n is 1 , section n is enabled.

### 3.3 Fixed-Point Facility Instructions

### 3.3.1 Fixed-Point Storage Access Instructions

The Storage Access instructions compute the effective address (EA) of the storage to be accessed as described in Section 1.11.3 on page 27.

## Programming Note

The la extended mnemonic permits computing an effective address as a Load or Store instruction would, but loads the address itself into a GPR rather than loading the value that is in storage at that address.

## Programming Note

The DS field in DS-form Storage Access instructions is a word offset, not a byte offset like the D field in D-form Storage Access instructions. However, for programming convenience, Assemblers should support the specification of byte offsets for both forms of instruction.

### 3.3.1.1 Storage Access Exceptions

Storage accesses will cause the system data storage error handler to be invoked if the program is not allowed to modify the target storage (Store only), or if the program attempts to access storage that is unavailable.

### 3.3.2 Fixed-Point Load Instructions

The byte, halfword, word, or doubleword in storage addressed by EA is loaded into register RT.
Many of the Load instructions have an "update" form, in which register RA is updated with the effective address. For these forms, if $R A \neq 0$ and $R A \neq R T$, the effective address is placed into register RA and the storage element (byte, halfword, word, or doubleword) addressed by EA is loaded into RT.

## Programming Note

In some implementations, the Load Algebraic and Load with Update instructions may have greater latency than other types of Load instructions. Moreover, Load with Update instructions may take longer to execute in some implementations than the corresponding pair of a non-update Load instruction and an Add instruction.

## Load Byte and Zero

D-form
lbz RT,D(RA)

| 34 | RT | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  | 3 |

```
if RA = 0 then b }\leftarrow
else b}\leftarrow(RA
EA}\leftarrow\textrm{b}+\operatorname{EXTS}(\textrm{D}
RT}\leftarrow\mp@subsup{}{}{56}0 || MEM(EA, 1
```

Let the effective address (EA) be the sum (RAIO)+ D. The byte in storage addressed by EA is loaded into $\mathrm{RT}_{56: 63} . \mathrm{RT}_{0: 55}$ are set to 0 .
Special Registers Altered:
None

## Load Byte and Zero with Update <br> D-form

Ibzu RT,D(RA)

| 35 | RT | RA |  | D | 31 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 | 16 |  | 3 |

```
EA \leftarrow (RA) + EXTS (D)
RT}\leftarrow\mp@subsup{}{}{56}0||\operatorname{MEM}(EA, 1
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ D. The byte in storage addressed by EA is loaded into $R T_{56: 63}$. $R T_{0: 55}$ are set to 0 .
EA is placed into register RA.
If $R A=0$ or $R A=R T$, the instruction form is invalid.

## Special Registers Altered:

None

Load Byte and Zero Indexed X-form
Ibzx RT,RA,RB

| 31 | RT | RA | RB |  | 87 | $/$ |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 | 16 | 21 |  |  |

```
if RA = 0 then b}\leftarrow
else }\quad\textrm{b}\leftarrow(RA
EA}\leftarrow\textrm{b}+(\textrm{RB}
RT}\leftarrow\mp@subsup{}{}{56}0||MEM(EA, 1
```

Let the effective address (EA) be the sum (RAIO)+ (RB). The byte in storage addressed by EA is loaded into $R T_{56: 63} . \mathrm{RT}_{0: 55}$ are set to 0 .
Special Registers Altered: None

## Load Byte and Zero with Update Indexed

X-form
Ibzux RT,RA,RB

| 31 | RT | RA | RB |  | 119 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

```
EA}\leftarrow(\textrm{RA})+(\textrm{RB}
RT}\leftarrow\mp@subsup{}{}{56}0||\operatorname{MEM}(EA,1
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ (RB). The byte in storage addressed by EA is loaded into $R T_{56: 63} . \mathrm{RT}_{0: 55}$ are set to 0 .
$E A$ is placed into register RA.
If $R A=0$ or $R A=R T$, the instruction form is invalid.
Special Registers Altered:
None

| Load Halfword and Zero |  |  |  | D-form |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| lhz RT,D(RA) |  |  |  |  |  |
| 40 | RT | RA |  | D |  |
|  | 6 | 11 | 16 |  | 31 |
| $\begin{aligned} & \text { if } R A=0 \text { then } b \leftarrow 0 \\ & \text { else } \quad b \leftarrow(R A) \end{aligned}$ |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
| EA $\leftarrow \mathrm{b}+\operatorname{EXTS}(\mathrm{D})$ <br> $\mathrm{RT} \leftarrow{ }^{48} 0\| \| \operatorname{MEM}(E A, 2)$ |  |  |  |  |  |
| Let the effective address (EA) be the sum (RAlO)+D. The halfword in storage addressed by EA is loaded into $\mathrm{RT}_{48: 63} \cdot \mathrm{RT}_{0: 47}$ are set to 0 . |  |  |  |  |  |
|  |  |  |  |  |  |
|  |  |  |  |  |  |
| Special Registers Altered: None |  |  |  |  |  |
| Load Halfword and Zero with Update D-form |  |  |  |  |  |
|  |  |  |  |  |  |
| Ihzu RT,D(RA) |  |  |  |  |  |
| 41 | RT | RA |  | D |  |
|  | 6 | 11 | 16 |  | 31 |
| $\mathrm{EA} \leftarrow(\mathrm{RA})+\operatorname{EXTS}(\mathrm{D})$ <br> $R T \leftarrow{ }^{48} 0 \\| \operatorname{MEM}(E A, 2)$ |  |  |  |  |  |
|  |  |  |  |  |  |
| $R A \leftarrow E A$ |  |  |  |  |  |
| Let the effective address (EA) be the sum (RA)+ D. The halfword in storage addressed by EA is loaded into $\mathrm{RT}_{48: 63} \cdot \mathrm{RT}_{0: 47}$ are set to 0 . |  |  |  |  |  |
| $E A$ is placed into register RA. |  |  |  |  |  |
| If $R A=0$ or $R A=R T$, the instruction form is invalid. |  |  |  |  |  |
| Special Registers Altered: None |  |  |  |  |  |

Load Halfword and Zero Indexed X-form
Ihzx RT,RA,RB

| 31 | RT | RA | RB |  | 279 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

```
if RA = 0 then b }\leftarrow
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
RT}\leftarrow\mp@subsup{}{}{480 || MEM(EA, 2)
```

Let the effective address (EA) be the sum (RAIO)+ (RB). The halfword in storage addressed by $E A$ is loaded into $\mathrm{RT}_{48: 63} . \mathrm{RT}_{0: 47}$ are set to 0 .

Special Registers Altered:
None

\section*{Load Halfword and Zero with Update Indexed <br> X-form <br> Ihzux RT,RA,RB <br> | 31 | RT | RA | RB |  | 311 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |}

```
EA}\leftarrow(\textrm{RA})+(\textrm{RB}
RT \leftarrow 480 || MEM(EA, 2)
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ (RB). The halfword in storage addressed by EA is loaded into $\mathrm{RT}_{48: 63} . \mathrm{RT}_{0: 47}$ are set to 0 .

EA is placed into register RA.
If $R A=0$ or $R A=R T$, the instruction form is invalid.
Special Registers Altered:
None

## Load Halfword Algebraic

Ina RT,D(RA)

| 42 | RT | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  |  |

Let the effective address (EA) be the sum (RAIO)+ D. The halfword in storage addressed by EA is loaded into $R T_{48: 63} . R T_{0: 47}$ are filled with a copy of bit 0 of the loaded halfword.

## Special Registers Altered:

None

## Load Halfword Algebraic with Update

 D-formIhau RT,D(RA)

| 43 | RT | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  |  |

```
EA \leftarrow(RA) + EXTS(D)
RT \leftarrow EXTS (MEM (EA, 2))
RA}\leftarrow\textrm{EA
\(\mathrm{EA} \leftarrow(\mathrm{RA})+\operatorname{EXTS}(\mathrm{D})\)
\(\mathrm{RT} \leftarrow \operatorname{EXTS}(\operatorname{MEM}(\operatorname{EA}, 2)\)
\(\mathrm{RA} \leftarrow \mathrm{EA}\)
```

Let the effective address (EA) be the sum (RA)+ D. The halfword in storage addressed by EA is loaded into $R T_{48: 63} \cdot R T_{0: 47}$ are filled with a copy of bit 0 of the loaded halfword. $E A$ is placed into register RA.
If $R A=0$ or $R A=R T$, the instruction form is invalid.

## Special Registers Altered:

None
D-form

$$
\square
$$

```
if RA = 0 then b }\leftarrow
```

if RA = 0 then b }\leftarrow
else b
else b
EA}\leftarrow\textrm{b}+\operatorname{EXTS}(\textrm{D}
EA}\leftarrow\textrm{b}+\operatorname{EXTS}(\textrm{D}
RT \leftarrow EXTS(MEM(EA, 2))

```
RT \leftarrow EXTS(MEM(EA, 2))
```

D-form

## Load Halfword Algebraic Indexed X-form

Ihax RT,RA,RB

| 31 | RT | RA | RB |  | 343 | $/$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

```
if RA = 0 then b }\leftarrow
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
RT \leftarrow EXTS(MEM(EA, 2))
```

Let the effective address (EA) be the sum (RAIO)+ (RB). The halfword in storage addressed by EA is loaded into $\mathrm{RT}_{48: 63} \cdot \mathrm{RT}_{0: 47}$ are filled with a copy of bit 0 of the loaded halfword.

Special Registers Altered:
None

## Load Halfword Algebraic with Update Indexed <br> X-form

Ihaux RT,RA,RB

| 31 | RT | RA | RB | 375 |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |

```
EA}\leftarrow(RA)+(RB
RT \leftarrow EXTS (MEM (EA, 2))
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ (RB). The halfword in storage addressed by EA is loaded into $\mathrm{RT}_{48: 63} . \mathrm{RT}_{0: 47}$ are filled with a copy of bit 0 of the loaded halfword.

EA is placed into register RA.
If $R A=0$ or $R A=R T$, the instruction form is invalid.

## Special Registers Altered:

 NoneLoad Word and Zero
if RA = 0 then b }\leftarrow
if RA = 0 then b }\leftarrow
else b b (RA)
else b b (RA)
EA \leftarrowb b EXTS(D)
EA \leftarrowb b EXTS(D)
RT}\leftarrow\mp@subsup{}{}{320 || MEM (EA, 4)
RT}\leftarrow\mp@subsup{}{}{320 || MEM (EA, 4)

D-form
Iwz RT,D(RA)

| 32 | RT | RA |  | D |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 | 16 |  |

Let the effective address (EA) be the sum (RAIO)+ D. The word in storage addressed by EA is loaded into $\mathrm{RT}_{32: 63} . \mathrm{RT}_{0: 31}$ are set to 0 .
Special Registers Altered:
None

## Load Word and Zero with Update D-form

## Iwzu RT,D(RA)

| 33 | RT | RA |  | D | 31 |
| :--- | :--- | :---: | :---: | :---: | ---: |
| 0 |  | 6 | 11 | 16 |  |

```
EA \leftarrow(RA) + EXTS(D)
RT}\leftarrow\mp@subsup{}{}{320}||MEM(EA,4
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+D. The word in storage addressed by EA is loaded into $\mathrm{RT}_{32: 63} . \mathrm{RT}_{0: 31}$ are set to 0 .

EA is placed into register RA.
If $R A=0$ or $R A=R T$, the instruction form is invalid.
Special Registers Altered:
None

Load Word and Zero Indexed X-form
Iwzx RT,RA,RB

| 31 | RT | RA | RB |  | 23 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

```
if RA = 0 then b }\leftarrow
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
RT \leftarrow 320 || MEM (EA, 4)
```

Let the effective address (EA) be the sum (RAIO)+ (RB). The word in storage addressed by EA is loaded into $\mathrm{RT}_{32: 63} \cdot \mathrm{RT}_{0: 31}$ are set to 0 .
Special Registers Altered:
None

Load Word and Zero with Update Indexed X-form

Iwzux RT,RA,RB

| 31 | RT | RA | RB |  | 55 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 |  | 21 |  |
| 16 |  |  |  |  |  |  |

```
EA}\leftarrow(\textrm{RA})+(\textrm{RB}
RT}\leftarrow\mp@subsup{}{}{32}0||\operatorname{MEM}(EA,4
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ (RB). The word in storage addressed by EA is loaded into $\mathrm{RT}_{32: 63} . \mathrm{RT}_{0: 31}$ are set to 0 .
EA is placed into register RA.
If $R A=0$ or $R A=R T$, the instruction form is invalid.
Special Registers Altered:
None

## Version 3.0

### 3.3.2.1 64-bit Fixed-Point Load Instructions



Let the effective address (EA) be the sum (RAIO)+ (DSIIOb00). The word in storage addressed by $E A$ is loaded into $R T_{32: 63} . R T_{0: 31}$ are filled with a copy of bit 0 of the loaded word.

## Special Registers Altered:

None

```
if RA = 0 then b }\leftarrow
be b & (RA)
EA \leftarrow b + EXTS(DS || 0b00)
RT \leftarrow EXTS(MEM(EA, 4))
```

```
if RA = 0 then b & 0
EA}\leftarrow\textrm{b}+(\textrm{RB}
RT \leftarrow EXTS(MEM(EA, 4))
```

Let the effective address (EA) be the sum (RAIO)+ (RB). The word in storage addressed by EA is loaded into $R T_{32: 63} . R T_{0: 31}$ are filled with a copy of bit 0 of the loaded word.

## Special Registers Altered:

None

## Load Word Algebraic with Update Indexed

 X-formIwaux RT,RA,RB

| 31 | RT | RA | RB |  | 373 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |

```
EA}\leftarrow(RA)+(RB
RT \leftarrow EXTS (MEM (EA, 4))
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ (RB). The word in storage addressed by EA is loaded into $R T_{32: 63} . R T_{0: 31}$ are filled with a copy of bit 0 of the loaded word.

EA is placed into register RA.
If $R A=0$ or $R A=R T$, the instruction form is invalid.
Special Registers Altered: None
Load Doubleword
Id
Id

| 58 | RT,DS(RA) |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  | 11 | RA |

```
if RA = 0 then b }\leftarrow
else b }\leftarrow(RA
EA \leftarrow b + EXTS(DS || Ob00)
RT \leftarrowMEM(EA, 8)
Let the effective address (EA) be the sum (RAIO)+ (DSIIOb00). The doubleword in storage addressed by EA is loaded into RT.
Special Registers Altered: None
```


## Load Doubleword with Update DS-form

Idu RT,DS(RA)

| 58 | RT | RA |  | DS | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 |  |

```
EA\leftarrow(RA) + EXTS(DS || 0b00)
RT \leftarrowMEM(EA, 8)
```

$\mathrm{RA} \leftarrow \mathrm{EA}$

Let the effective address (EA) be the sum (RA)+ (DSIIObOO). The doubleword in storage addressed by EA is loaded into RT.

EA is placed into register RA.
If $R A=0$ or $R A=R T$, the instruction form is invalid.
Special Registers Altered:
None


```
if RA = 0 then b }\leftarrow
else b 
EA}\leftarrow\textrm{b}+(\textrm{RB}
RT \leftarrow MEM(EA, 8)
Let the effective address (EA) be the sum
(RAIO)+ (RB). The doubleword in storage addressed by
EA is loaded into RT.
Special Registers Altered:
    None
```


## Load Doubleword with Update Indexed

 $X$-formIdux RT,RA,RB

| 31 | RT | RA | RB |  | 53 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

```
EA}\leftarrow(RA)+(RB
RT}\leftarrowMEM(EA, 8
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ (RB). The doubleword in storage addressed by EA is loaded into RT.

EA is placed into register RA.
If $R A=0$ or $R A=R T$, the instruction form is invalid.
Special Registers Altered:
None

## Load Doubleword Monitored Indexed X-form

| Idmx RT,RA,RB |
| :--- |
| 31 RT RA RB 309 $/$ <br> 0 6  11 16 21 |

```
if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(RA
EA \leftarrow b + (RB)
loaded_ea }\leftarrow\operatorname{MEM(EA, 8)
if ᄀ((loaded_ea is in enabled section of
    load-monitored region) & BESCR (GE LME}=0b11
RT \leftarrow loaded_ea
```

Let the effective address (EA) be the sum (RAIO) $+(R B)$.
The doubleword in storage addressed by EA is loaded. If the loaded doubleword, interpreted as an effective address, is not equal to an effective address within an enabled section of the Load Monitored region, or if Load Monitored event-based branches are disabled ( BESCR $_{\text {LME }}=0$ or $B_{E S C R}^{G E}=0$ ), the loaded doubleword is placed into RT; otherwise a Load Monitored event-based branch occurs. (The Load Monitored region is described in Section 3.2.4. Event-based branches and the BESCR are described in Chapter 7 of Book II.)
The doubleword in storage specified by the Idmx instruction should not be in storage that is both Caching Inhibited and Guarded; execution of an Idmx instruction to access storage with these attributes will result in boundedly undefined behavior. (See Section 1.6.2 of Book II and Section 1.6.4 of Book II.

In privileged state, this instruction is an illegal instruction and an attempt to execute it will invoke the system illegal instruction error handler. See Section 4.4.3 of Book III for additional information.

## Special Registers Altered:

None

## Programming Note

Warning: The Idmx instruction should be in storage to which the event-based branch handler has read access (see Section 5.7.14 of Book III), because the event-based branch handler may need to load the Idmx instruction from storage in order to determine the instruction's register usage.

## Programming Note

The Idmx instruction loads the same doubleword in storage as the Idx instruction.

## - Programming Note

Idmx is intended for use by applications when loading data objects during times when objects are being moved as part of a garbage collection process. In this type of usage, all loads of object pointers by applications are performed using Idmx.

Whenever objects within a given region of memory are to be moved, the garbage collection program sets the Load Monitored Region registers to encompass the region being moved, and sets $\mathrm{BESCR}_{G E}$ LME to Ob11 to enable Load Monitored event-based branches.
Subsequently, if an application program loads (Idmx) a pointer into an enabled section of the Load Monitored region, an event-based branch (EBB) will occur. The EBB handler will load the instruction from storage, and decode the instruction to determine the effective address from which the pointer was loaded. The handler will then load the pointer (obtaining the same value the application obtained), and determine where the corresponding object has been moved to, and update the pointer in storage so that it points to the object's new location. The EBB handler may also take other actions such as updating additional pointers, depending on the situation. After this processing is complete, the EBB handler sets BESCR LME LMEO to (10) since the taking of the EBB set these bits to $(0,1)$, and then executes rfebb 1 to re-enable EBBs and return to the application program at the Idmx. The application program re-executes the Idmx, which now returns the updated pointer (which no longer points into an enabled section of the Load Monitored region), and continues.

Other variations and extensions of the above procedure are also possible.

### 3.3.3 Fixed-Point Store Instructions

The contents of register RS are stored into the byte, halfword, word, or doubleword in storage addressed by EA.

Many of the Store instructions have an "update" form, in which register RA is updated with the effective address. For these forms, the following rules apply.

- If $R A \neq 0$, the effective address is placed into register RA.
- If $R S=R A$, the contents of register RS are copied to the target storage element and then EA is placed into RA (RS).

| Store Byte |  |  |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| stb |  |  |  |  |  |  |
| RS,D(RA) |  |  |  |  |  | D-form |
| 38  RS RA  D |  |  |  |  |  |  |
| 0 |  |  |  |  |  |  |

```
if RA = 0 then b b 
else b}\leftarrow(RA
EA \leftarrow b + EXTS(D)
MEM(EA, 1) \leftarrow(RS) 56:63
```

Let the effective address (EA) be the sum (RAIO)+ $D$. $(\mathrm{RS})_{56: 63}$ are stored into the byte in storage addressed by EA.

## Special Registers Altered:

None
Store Byte with Update
stbu

| 39 | $R S, D(R A)$ | D-form |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 |  | 11 | 16 | DA |  |

```
EA \leftarrow(RA) + EXTS(D)
MEM(EA, 1) \leftarrow(RS) 56:63
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+D. $(\mathrm{RS})_{56: 63}$ are stored into the byte in storage addressed by EA.

EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None
Store Byte Indexed
stbx

| 31 | RS,RA,RB |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 |  | RS | RA | RB |

if $\mathrm{RA}=0$ then $\mathrm{b} \leftarrow 0$

```
else b}\leftarrow(RA
```

$\mathrm{EA} \leftarrow \mathrm{b}+(\mathrm{RB})$
$\operatorname{MEM}(E A, 1) \leftarrow(R S)_{56: 63}$

Let the effective address (EA) be the sum (RAIO)+ (RB). (RS) $)_{56: 63}$ are stored into the byte in storage addressed by EA.

## Special Registers Altered:

None

## Store Byte with Update Indexed X-form

stbux RS,RA,RB

| 31 | RS | RA | RB |  | 247 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 1 |  |  |  |  |  |  |

```
EA}\leftarrow(RA)+(RB
MEM(EA, 1) \leftarrow(RS) 56:63
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ (RB). $(\mathrm{RS})_{56: 63}$ are stored into the byte in storage addressed by EA.

EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None

## Store Halfword

sth RS,D(RA)

| 44 | RS | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  |  |

```
```

if RA = 0 then b }\leftarrow

```
```

if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(\textrm{RA}
else }\quad\textrm{b}\leftarrow(\textrm{RA}
EA \leftarrow b + EXTS(D)
EA \leftarrow b + EXTS(D)
MEM(EA, 2) \leftarrow (RS) 48:63

```
```

MEM(EA, 2) \leftarrow (RS) 48:63

```
```

Let the effective address (EA) be the sum (RAIO)+ D. $(\mathrm{RS})_{48: 63}$ are stored into the halfword in storage addressed by EA.

## Special Registers Altered:

None

Store Halfword with Update
D-form
sthu RS,D(RA)

| 45 | RS | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  | 3 |

```
EA \leftarrow(RA) + EXTS (D)
MEM(EA, 2) \leftarrow(RS) 48:63
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ D. $(\mathrm{RS})_{48: 63}$ are stored into the halfword in storage addressed by EA.
EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None
D-form
-

31

Store Halfword with Update Indexed X-form
sthux RS,RA,RB

| 31 | RS | RA | RB |  | 439 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

```
EA \leftarrow (RA) + (RB)
MEM(EA, 2) \leftarrow(RS) 48:63
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ (RB). $(\mathrm{RS})_{48: 63}$ are stored into the halfword in storage addressed by EA.
$E A$ is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None


```
if RA = 0 then b }\leftarrow
else b
EA \leftarrow b + EXTS (D)
MEM(EA, 4)\leftarrow(RS) 32:63
```

Let the effective address (EA) be the sum (RAIO)+ D. (RS) ${ }_{32: 63}$ are stored into the word in storage addressed by EA.
Special Registers Altered:
None

## Store Word with Update

D-form
stwu RS,D(RA)

| 37 | RS | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |

```
EA \leftarrow(RA) + EXTS(D)
MEM(EA,4) 4 (RS) 32:63
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+ D. $(\mathrm{RS})_{32: 63}$ are stored into the word in storage addressed by EA.

EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None

Store Word Indexed X-form
stwx RS,RA,RB

| 31 | RS | RA | RB |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 151 | $/$ |  |

```
if RA = 0 then b \leftarrow 0
else b b & RA)
EA}\leftarrow\textrm{b}+(\textrm{RB}
MEM(EA, 4) \leftarrow(RS) 32:63
```

Let the effective address (EA) be the sum (RAIO)+ (RB). (RS) ${ }_{32: 63}$ are stored into the word in storage addressed by EA.

Special Registers Altered:
None

## Store Word with Update Indexed X-form

stwux RS,RA,RB

| 31 | RS | RA | RB |  | 183 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

```
EA \leftarrow (RA) + (RB)
MEM(EA,4) \leftarrow(RS) 32:63
RA}\leftarrowE
```

Let the effective address (EA) be the sum (RA)+ (RB). $(\mathrm{RS})_{32: 63}$ are stored into the word in storage addressed by EA.

EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None

### 3.3.3.1 64-bit Fixed-Point Store Instructions

Store Doubleword
std

| 62 | RS | RS-form |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | DS |

```
if RA = 0 then b }\leftarrow
else b
EA \leftarrow b + EXTS(DS || 0b00)
MEM(EA, 8) \leftarrow (RS)
```

Let the effective address (EA) be the sum (RAIO)+ (DSIIOb00). (RS) is stored into the doubleword in storage addressed by EA.

## Special Registers Altered: None <br> Store Doubleword with Update <br> DS-form

stdu RS,DS(RA)

| 62 | RS | RA |  | DS | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 |  |

```
EA\leftarrow(RA) + EXTS(DS || 0b00)
MEM(EA, 8) \leftarrow(RS)
RA}\leftarrowE
Let the effective address (EA) be the sum (RA)+ (DSIIOb00). (RS) is stored into the doubleword in
``` storage addressed by EA.
\(E A\) is placed into register RA.
If \(R A=0\), the instruction form is invalid.
Special Registers Altered:
None

Store Doubleword Indexed X-form
stdx RS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 149 \\
\hline 0 & & & 11 & 16 & 21 \\
& & & \\
\hline
\end{tabular}
```

if RA = 0 then b \leftarrow }\leftarrow
else }\quad\textrm{b}\leftarrow(\textrm{RA}
EA}\leftarrow\textrm{b}+(\textrm{RB}

```
\(\operatorname{MEM}(E A, 8) \leftarrow(R S)\)

Let the effective address (EA) be the sum (RAIO)+ (RB). (RS) is stored into the doubleword in storage addressed by EA.

\section*{Special Registers Altered:} None

Store Doubleword with Update Indexed \(X\)-form
stdux RS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 181 & 1 \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

EA}\leftarrow(\textrm{RA})+(\textrm{RB}
MEM (EA, 8) \leftarrow(RS)
RA}\leftarrowE

```

Let the effective address (EA) be the sum (RA)+ (RB). (RS) is stored into the doubleword in storage addressed by EA.
EA is placed into register RA.
If \(R A=0\), the instruction form is invalid.
Special Registers Altered:
None

\subsection*{3.3.4 Fixed Point Load and Store Quadword Instructions}

For Iq, the quadword in storage addressed by EA is loaded into an even-odd pair of GPRs as follows. In Big-Endian mode, the even-numbered GPR is loaded with the doubleword from storage addressed by EA and the odd-numbered GPR is loaded with the doubleword addressed by EA+8. In Little-Endian mode, the even-numbered GPR is loaded with the byte-reversed doubleword from storage addressed by EA+8 and the odd-numbered GPR is loaded with the byte-reversed doubleword addressed by EA.

In the preferred form of the Load Qudword instruction \(R A \neq R T p+1\).
For \(\boldsymbol{s t q}\), the contents of an even-odd pair of GPRs is stored into the quadword in storage addressed by EA as follows. In Big-Endian mode, the even-numbered GPR is stored into the doubleword in storage addressed by EA and the odd-numbered GPR is stored into the doubleword addressed by EA+8. In Lit-tle-Endian mode, the even-numbered GPR is stored byte-reversed into the doubleword in storage addressed by EA+8 and the odd-numbered GPR is stored byte-reversed into the doubleword addressed by EA.

\section*{Programming Note}

The Iq and stq instructions exist primarily to permit software to access quadwords in storage "atomically"; see Section 1.4 of Book II. Because GPRs are 64 bits long, the Fixed-Point Facility on many designs is optimized for storage accesses of at most eight bytes. On such designs, the quadword atomicity required for \(\boldsymbol{I q}\) and \(\boldsymbol{s t q}\) makes these instructions complex to implement, with the result that the instructions may perform less well on these designs than the corresponding two Load Doubleword or Store Doubleword instructions.

The complexity of providing quadword atomicity may be especially great for storage that is Write Through Required or Caching Inhibited (see Section 1.6 of Book II). This is why Iq and stq are permitted to cause the data storage error handler to be invoked if the specified storage location is in either of these kinds of storage (see Section 3.3.1.1).

\section*{Load Quadword}
\(D Q\)-form
lq \(\quad \mathrm{RTp}, \mathrm{DQ}(\mathrm{RA})\)
\begin{tabular}{|c|c|c|cc|c|}
\hline 56 & RTp & RA & & DQ & \begin{tabular}{c} 
I/I \\
0
\end{tabular} \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else b
EA \leftarrow b + EXTS(DQ || 0b0000)
RTp \leftarrowMEM(EA, 16)

```

Let the effective address (EA) be the sum (RAIO)+ (DQlIOb0000). The quadword in storage addressed by EA is loaded into register pair RTp.

If RTp is odd or RTp=RA, the instruction form is invalid. If \(R T p=R A\), an attempt to execute this instruction will invoke the system illegal instruction error handler. (The \(R T p=R A\) case includes the case of \(R T p=R A=0\).)
The quadword in storage addressed by EA is loaded into an even-odd pair of GPRs as follows. In Big-Endian mode, the even-numbered GPR is loaded with the doubleword from storage addressed by EA and the odd-numbered GPR is loaded with the doubleword addressed by EA+8. In Little-Endian mode, the even-numbered GPR is loaded with the byte-reversed doubleword from storage addressed by EA+8 and the odd-numbered GPR is loaded with the byte-reversed doubleword addressed by EA.

\section*{Programming Note}

In versions of the architecture prior to V. 2.07, this instruction was privileged.

\section*{Special Registers Altered:}

None

\section*{Version 3.0}

\section*{Store Quadword}

DS-form
stq \(\quad\) RSp,DS(RA)
\begin{tabular}{|c|c|c|cc|c|}
\hline 62 & RSp & RA & & DS & 2 \\
0 & & 6 & 11 & 16 & \\
3031 \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else b}\leftarrow(RA
EA \leftarrow b + EXTS(DS || 0b00)
MEM(EA, 16) \leftarrowRSp

```

Let the effective address (EA) be the sum (RAIO)+ (DSIIObOO). The contents of register pair RSp are stored into the quadword in storage addressed by EA.

If RSp is odd, the instruction form is invalid.
The contents of an even-odd pair of GPRs is stored into the quadword in storage addressed by EA as follows. In Big-Endian mode, the even-numbered GPR is stored into the doubleword in storage addressed by EA and the odd-numbered GPR is stored into the doubleword addressed by EA+8. In Little-Endian mode, the even-numbered GPR is stored byte-reversed into the doubleword in storage addressed by EA+8 and the odd-numbered GPR is stored byte-reversed into the doubleword addressed by EA.

\section*{Programming Note}

In versions of the architecture prior to V. 2.07, this instruction was privileged.

Special Registers Altered:
None

\subsection*{3.3.5 Fixed-Point Load and Store with Byte Reversal Instructions}

\section*{Programming Note}

These instructions have the effect of loading and storing data in the opposite byte ordering from that which would be used by other Load and Store instructions.

\section*{Programming Note}

In some implementations, the Load Byte-Reverse instructions may have greater latency than other Load instructions.

\section*{Load Halfword Byte-Reverse Indexed X-form}
Ihbrx \begin{tabular}{l} 
RT,RA,RB \\
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & & 790 & 1 \\
0 & & 6 & & 11 & 16 & 21 \\
31 \\
\hline
\end{tabular}
\end{tabular}\(.\)\begin{tabular}{l} 
\\
\hline
\end{tabular}
```

if $\mathrm{RA}=0$ then $\mathrm{b} \leftarrow 0$
else $\quad b \leftarrow(R A)$
$\mathrm{EA} \leftarrow \mathrm{b}+(\mathrm{RB})$
load_data $\leftarrow$ MEM $(E A, 2)$
$R T \leftarrow{ }^{48} 0| |$ load_data $_{8: 15}| |$ load_data $_{0: 7}$

```

Let the effective address (EA) be the sum (RAIO)+(RB). Bits 0:7 of the halfword in storage addressed by EA are loaded into \(\mathrm{RT}_{56: 63}\). Bits 8:15 of the halfword in storage addressed by EA are loaded into \(\mathrm{RT}_{48: 55} . \mathrm{RT}_{0: 47}\) are set to 0 .

\section*{Special Registers Altered:}

None

\section*{Load Word Byte-Reverse Indexed X-form}
Iwbrx RT,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & & 534 & 1 \\
\hline 0 & & & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else b}\leftarrow(RA
EA \leftarrow b + (RB)
load_data \leftarrow MEM(EA, 4)
RT}\leftarrow\mp@subsup{}{}{320}|| load_data 24:31 || load_data 16:23
|| load_data_8:15 || load_data0:7

```

Let the effective address (EA) be the sum (RAIO)+ (RB). Bits 0:7 of the word in storage addressed by EA are loaded into \(\mathrm{RT}_{56: 63}\). Bits 8:15 of the word in storage addressed by EA are loaded into \(\mathrm{RT}_{48: 55}\). Bits 16:23 of the word in storage addressed by EA are loaded into \(\mathrm{RT}_{40: 47}\). Bits 24:31 of the word in storage addressed by EA are loaded into \(\mathrm{RT}_{32: 39}\). \(\mathrm{RT}_{0: 31}\) are set to 0 .

\section*{Special Registers Altered:}

None

\section*{Store Halfword Byte-Reverse Indexed X-form}
sthbrx RS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 918 & 1 \\
\hline 0 & & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
MEM (EA, 2) \leftarrow (RS) 56:63 || (RS) 48:55

```

Let the effective address (EA) be the sum (RAIO) \(+(\mathrm{RB}) .(\mathrm{RS})_{56: 63}\) are stored into bits 0:7 of the halfword in storage addressed by EA. (RS) 48:55 \(^{\text {are }}\) stored into bits \(8: 15\) of the halfword in storage addressed by EA.

\section*{Special Registers Altered: \\ None}

\section*{Store Word Byte-Reverse Indexed X-form}
stwbrx RS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 662 & 1 \\
\hline 0 & & & & 11 & 16 & 21 \\
& & & 31 \\
\hline
\end{tabular}
```

if RA = 0 then b b
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
MEM(EA,4)\leftarrow(RS) 56:63 || (RS) 48:55 || (RS) 40:47
| ( (RS) 32:39

```

Let the effective address (EA) be the sum (RAIO)+ (RB). (RS \()_{56: 63}\) are stored into bits 0:7 of the word in storage addressed by EA. (RS) \(48: 55\) are stored into bits \(8: 15\) of the word in storage addressed by EA. \((R S)_{40: 47}\) are stored into bits 16:23 of the word in storage addressed by EA. (RS) 32:39 \(^{2}\) are stored into bits 24:31 of the word in storage addressed by EA.

\section*{Special Registers Altered:}

None

\subsection*{3.3.5.1 64-Bit Load and Store with Byte Reversal Instructions}

\section*{Load Doubleword Byte-Reverse Indexed \\ X-form}

Idbrx
RT,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & & RT & RA & RB & & 532 \\
\hline 0 & & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(\textrm{RA
EA}\leftarrow\textrm{b}+(\textrm{RB}
load_data \leftarrow MEM(EA, 8)
RT \leftarrow load_data56:63 || load_data 48:55
|| load_data40:47 || load_data 32:39
| load_data8:15 || load_data0:7

```

Let the effective address (EA) be the sum (RAIO)+(RB). Bits 0:7 of the doubleword in storage addressed by EA are loaded into \(\mathrm{RT}_{56: 63}\). Bits 8:15 of the doubleword in storage addressed by EA are loaded into \(\mathrm{RT}_{48: 55}\). Bits 16:23 of the doubleword in storage addressed by EA are loaded into \(\mathrm{RT}_{40: 47}\). Bits 24:31 of the doubleword in storage addressed by EA are loaded into \(\mathrm{RT}_{32: 39}\). Bits 32:39 of the doubleword in storage addressed by EA are loaded into \(\mathrm{RT}_{24: 31}\). Bits 40:47 of the doubleword in storage addressed by EA are loaded into \(\mathrm{RT}_{16: 23}\). Bits 48:55 of the doubleword in storage addressed by EA are loaded into \(R T_{8: 15}\). Bits 56:63 of the doubleword in storage addressed by EA are loaded into \(\mathrm{RT}_{0: 7}\).
Special Registers Altered:
None

\section*{Store Doubleword Byte-Reverse Indexed X-form}
stdbrx RS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{2}{|c|}{ RS } & RA & RB & & 660 \\
\hline 0 & & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

if RA = 0 then b \leftarrow 0
else b}
EA \leftarrow b + (RB)
MEM(EA, 8) \leftarrow(RS) 56:63 || (RS) 48:55

```


Let the effective address (EA) be the sum (RAIO)+(RB). (RS) \(56: 63\) are stored into bits \(0: 7\) of the doubleword in storage addressed by EA. (RS) 48:55 \(^{2}\) are stored into bits 8:15 of the doubleword in storage addressed by EA. (RS) 40:47 are stored into bits \(16: 23\) of the doubleword in storage addressed by EA. (RS) 32:39 \(^{2}\) are stored into bits 23:31 of the doubleword in storage addressed by EA. (RS) 24:31 are stored into bits 32:39 of the doubleword in storage addressed by EA. (RS) \(16: 23\) are stored into bits 40:47 of the doubleword in storage addressed by EA. (RS) \()_{8: 15}\) are stored into bits \(48: 55\) of the doubleword in storage addressed by EA. (RS) \({ }_{0: 7}\) are stored into bits 56:63 of the doubleword in storage addressed by EA.

\section*{Special Registers Altered:}

None

\subsection*{3.3.6 Fixed-Point Load and Store Multiple Instructions}

Load Multiple Word
Imw RT,D(RA)
\begin{tabular}{|l|l|l|lll|}
\hline 46 & RT & RA & \multicolumn{2}{c|}{} & D \\
0 & & 6 & & 11 & 16 \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(\textrm{RA}
EA \leftarrow b + EXTS(D)
r}\leftarrow\textrm{RT
do while r \leq 31
GPR(r)\leftarrow < 320 || MEM(EA, 4)
r}\leftarrowr+
EA}\leftarrow\textrm{EA}+

```

Let \(\mathrm{n}=(32-\mathrm{RT})\). Let the effective address (EA) be the sum (RAIO)+ D.
n consecutive words starting at EA are loaded into the low-order 32 bits of GPRs RT through 31. The high-order 32 bits of these GPRs are set to zero.

If RA is in the range of registers to be loaded, including the case in which \(R A=0\), the instruction form is invalid.

This instruction is not supported in Little-Endian mode. If it is executed in Little-Endian mode, the system alignment error handler is invoked.

Special Registers Altered:
None

Store Multiple Word
D-form
stmw RS,D(RA)
\begin{tabular}{|l|l|l|lll|}
\hline 47 & RS & RA & & D & 31 \\
\hline 0 & 6 & 11 & 16 & & 3 \\
\hline
\end{tabular}
if \(R A=0\) then \(b \leftarrow 0\)
```

else b}\leftarrow(RA
EA \leftarrow b + EXTS(D)
r}\leftarrow\textrm{RS
do while r s 31
MEM(EA, 4) \leftarrowGPR(r)}32:6
r}\leftarrowr+
EA \leftarrowEA + 4

```

Let \(\mathrm{n}=(32-\mathrm{RS})\). Let the effective address (EA) be the sum (RAl0)+ D.
n consecutive words starting at EA are stored from the low-order 32 bits of GPRs RS through 31.

This instruction is not supported in Little-Endian mode. If it is executed in Little-Endian mode, the system alignment error handler is invoked.

\section*{Special Registers Altered:}

None

\subsection*{3.3.7 Fixed-Point Move Assist Instructions [Phased Out]}

The Move Assist instructions allow movement of an arbitrary sequence of bytes from storage to registers or from registers to storage without concern for alignment. These instructions can be used for a short move between arbitrary storage locations or to initiate a long move between unaligned storage fields.

The Move Assist instructions have preferred forms; see Section 1.9.1, "Preferred Instruction Forms" on page 23. In the preferred forms, register usage satisfies the following rules.
- \(\mathrm{RS}=4\) or 5
- \(\mathrm{RT}=4\) or 5
- last register loaded/stored \(\leq 12\)

For some implementations, using GPR 4 for RS and RT may result in slightly faster execution than using GPR 5.

\section*{Load String Word Immediate X-form \\ Iswi RT,RA,NB \\ \begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & NB & & 597 & 1 \\
\hline 0 & & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}}
```

if RA = 0 then EA }\leftarrow
else EA }\leftarrow(RA
if NB = 0 then n }\leftarrow3
else }n<N
r}\leftarrowRT - 1
i}\leftarrow3
do while n > 0
if i = 32 then
r}\leftarrowr+1(\operatorname{mod}32
GPR(r) \leftarrow0
GPR(r) i:i+7}< \leftarrowMEM(EA, 1
i}\leftarrowi+
if i = 64 then i }\leftarrow3
EA}\leftarrow\textrm{EA}+
n}\leftarrow\textrm{n}-

```

Let the effective address (EA) be (RAIO). Let \(n=N B\) if \(N B \neq 0, n=32\) if \(N B=0 ; n\) is the number of bytes to load. Let \(n r=\operatorname{CEIL}(\mathrm{n} / 4)\); nr is the number of registers to receive data.
n consecutive bytes starting at EA are loaded into GPRs RT through RT+nr-1. Data are loaded into the low-order four bytes of each GPR; the high-order four bytes are set to 0 .

Bytes are loaded left to right in each register. The sequence of registers wraps around to GPR 0 if required. If the low-order four bytes of register RT+nr-1 are only partially filled, the unfilled low-order byte(s) of that register are set to 0 .

If RA is in the range of registers to be loaded, including the case in which \(R A=0\), the instruction form is invalid.

This instruction is not supported in Little-Endian mode. If it is executed in Little-Endian mode, the system alignment error handler is invoked.

\section*{Special Registers Altered:}

None

Load String Word Indexed X-form Iswx RT,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & & 533 & 1 \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

if $R A=0$ then $b \leftarrow 0$
else $\quad \mathrm{b} \leftarrow(R A)$
$\mathrm{EA} \leftarrow \mathrm{b}+(\mathrm{RB})$
$\mathrm{n} \leftarrow \mathrm{XER}_{57: 63}$
$r \leftarrow R T-1$
$i \leftarrow 32$
$\mathrm{RT} \leftarrow$ undefined
do while $\mathrm{n}>0$
if $i=32$ then
$r \leftarrow r+1(\bmod 32)$
$\operatorname{GPR}(r) \leftarrow 0$
$\operatorname{GPR}(r)_{i: i+7} \leftarrow \operatorname{MEM}(E A, 1)$
$i \leftarrow i+8$
if $i=64$ then $i \leftarrow 32$
$\mathrm{EA} \leftarrow \mathrm{EA}+1$
$\mathrm{n} \leftarrow \mathrm{n}-1$

```

Let the effective address (EA) be the sum (RAIO)+ (RB). Let \(n=X^{2} R_{57: 63} ; n\) is the number of bytes to load. Let \(\mathrm{nr}=\mathrm{CEIL}(\mathrm{n} / 4)\); nr is the number of registers to receive data.

If \(n>0, n\) consecutive bytes starting at EA are loaded into GPRs RT through RT+nr-1. Data are loaded into the low-order four bytes of each GPR; the high-order four bytes are set to 0 .

Bytes are loaded left to right in each register. The sequence of registers wraps around to GPR 0 if required. If the low-order four bytes of register RT+nr-1 are only partially filled, the unfilled low-order byte(s) of that register are set to 0 .
If \(n=0\), the contents of register RT are undefined.
If RA or RB is in the range of registers to be loaded, including the case in which \(R A=0\), the instruction is treated as if the instruction form were invalid. If \(R T=R A\) or \(R T=R B\), the instruction form is invalid.

This instruction is not supported in Little-Endian mode. If it is executed in Little-Endian mode and \(n>0\), the system alignment error handler is invoked.

\section*{Special Registers Altered:}

None

\section*{Store String Word Immediate}

X-form
stswi RS,RA,NB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & NB & & 725 & 1 \\
0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

if RA = 0 then EA \leftarrow0
else EA }\leftarrow(RA
if NB = 0 then n}\leftarrow3
else }\quad\textrm{n}\leftarrowN\textrm{N
r}\leftarrow\textrm{RS}-
i}\leftarrow3
do while n > 0
if i = 32 then r tr + 1 (mod 32)
MEM(EA, 1)}\leftarrow\operatorname{GPR}(r)\mp@subsup{)}{i:i+7}{
i}\leftarrowi+
if i = 64 then i }\leftarrow3
EA}\leftarrow\textrm{EA}+
n}\leftarrow\textrm{n}-

```

Let the effective address (EA) be (RAIO). Let \(\mathrm{n}=\mathrm{NB}\) if \(N B \neq 0, n=32\) if \(N B=0 ; n\) is the number of bytes to store. Let \(n r=C E I L(n / 4)\); \(n r\) is the number of registers to supply data.
n consecutive bytes starting at EA are stored from GPRs RS through RS+nr-1. Data are stored from the low-order four bytes of each GPR.

Bytes are stored left to right from each register. The sequence of registers wraps around to GPR 0 if required.
This instruction is not supported in Little-Endian mode. If it is executed in Little-Endian mode, the system alignment error handler is invoked.
Special Registers Altered: None

Store String Word Indexed
X-form
stswx RS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 661 & \(/\) \\
\hline 0 & & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
n}\leftarrow\mp@subsup{\textrm{XER}}{57:63}{
r}\leftarrowRS - 1
i}\leftarrow3
do while n > 0
if i = 32 then r tr + 1 (mod 32)
MEM(EA, 1)}\leftarrow\operatorname{GPR}(r)\mp@subsup{)}{i:i+7}{i
i}\leftarrowi+
if i = 64 then i }\leftarrow3
EA}\leftarrow\textrm{EA}+
n}\leftarrow\textrm{n}-

```

Let the effective address (EA) be the sum (RAIO)+ (RB). Let \(n=X_{5 E R} 57\); \(n\) is the number of bytes to store. Let \(\mathrm{nr}=\operatorname{CEIL}(\mathrm{n} / 4)\); nr is the number of registers to supply data.

If \(n>0, n\) consecutive bytes starting at EA are stored from GPRs RS through RS+nr-1. Data are stored from the low-order four bytes of each GPR.

Bytes are stored left to right from each register. The sequence of registers wraps around to GPR 0 if required.
If \(\mathrm{n}=0\), no bytes are stored.
This instruction is not supported in Little-Endian mode. If it is executed in Little-Endian mode and \(n>0\), the system alignment error handler is invoked.

\section*{Special Registers Altered:} None

\subsection*{3.3.8 Other Fixed-Point Instructions}

The remainder of the fixed-point instructions use the contents of the General Purpose Registers (GPRs) as source operands, and place results into GPRs, into the Fixed-Point Exception Register (XER), and into Condition Register fields. In addition, the Trap instructions test the contents of a GPR or XER bit, invoking the system trap handler if the result of the specified test is true.

These instructions treat the source operands as signed integers unless the instruction is explicitly identified as performing an unsigned operation.

The X-form and XO-form instructions with Rc=1, and the D-form instructions addic., andi., and andis., set the first three bits of CR Field 0 to characterize the result placed into the target register. In 64-bit mode,
these bits are set by signed comparison of the result to zero. In 32-bit mode, these bits are set by signed comparison of the low-order 32 bits of the result to zero.

Unless otherwise noted and when appropriate, when CR Field 0 and the XER are set they reflect the value placed into the target register.

\section*{Programming Note}

Instructions with the OE bit set or that set CA and CA32 may execute slowly or may prevent the execution of subsequent instructions until the instruction has completed.

\subsection*{3.3.9 Fixed-Point Arithmetic Instructions}

The XO-form Arithmetic instructions with Rc=1, and the D-form Arithmetic instruction addic., set the first three bits of CR Field 0 as described in Section 3.3.8, "Other Fixed-Point Instructions".
addic, addic., subfic, addc, subfc, adde, subfe, addme, subfme, addze, and subfze always set CA, to reflect the carry out of bit 0 in 64-bit mode and out of bit 32 in 32-bit mode. These instructions also always set CA32 to reflect the carry out of bit 32. The XO-form Arithmetic instructions set SO, OV, and OV32 when \(\mathrm{OE}=1\) to reflect overflow of the result. Except for the Multiply Low and Divide instructions, the setting of SO and OV bits is mode-dependent, and reflects overflow of the 64-bit result in 64-bit mode and overflow of the low-order 32-bit result in 32-bit mode, while OV32 reflects overflow of the low-order 32-bit result independent of the mode. For XO-form Multiply Low and Divide instructions, the setting of SO, OV, and OV32 bits is mode-independent, and reflects overflow of the 64-bit result for mulld, divd, divde, divdu and divdeu, and overflow of the low-order 32-bit result for mullw, divw, divwe, divwu, and divweu.

\section*{Programming Note}

Notice that CR Field 0 may not reflect the "true" (infinitely precise) result if overflow occurs.

\section*{Extended mnemonics for addition and subtraction}

Several extended mnemonics are provided that use the Add Immediate and Add Immediate Shifted instructions to load an immediate value or an address into a target register. Some of these are shown as examples with the two instructions.

The Power ISA supplies Subtract From instructions, which subtract the second operand from the third. A set of extended mnemonics is provided that use the more "normal" order, in which the third operand is subtracted from the second, with the third operand being either an immediate field or a register. Some of these are shown as examples with the appropriate Add and Subtract From instructions.

See Appendix C for additional extended mnemonics.
Add Immediate
addi
\begin{tabular}{|l|l|l|lll|}
\hline 14 & RT,RA,SI & D-form \\
\hline 0 & 6 & 11 & 16 & SI & \\
\hline
\end{tabular}
```

if RA = 0 then RT }\leftarrow\operatorname{EXTS}(SI
else RT \leftarrow (RA) + EXTS(SI)

```

The sum (RAl0) + SI is placed into register RT.

\section*{Special Registers Altered:} None

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Add Immediate:
\begin{tabular}{ll}
\multicolumn{2}{l}{ Extended: } \\
li & \(R x\), value \\
la & \(R x\), disp(Ry) \\
subi & \(R x, R y\), value
\end{tabular}

\section*{Equivalent to:}
addi \(R x, 0\),value
addi \(R x, R y, d i s p\)
addi \(R x, R y\),-value

\section*{Programming Note}
addi, addis, add, and subf are the preferred instructions for addition and subtraction, because they set few status bits.

Notice that addi and addis use the value 0, not the contents of GPR 0 , if RA=0.

Add Immediate Shifted D-form
addis RT,RA,SI
\begin{tabular}{|l|l|l|lll|}
\hline 15 & RT & RA & \multicolumn{2}{|c|}{} & SI \\
\hline 0 & 6 & & 11 & 16 & \\
\hline
\end{tabular}
if \(R A=0\) then \(R T \leftarrow \operatorname{EXTS}\left(S I\left|\mid{ }^{16} 0\right)\right.\)
else \(R T \leftarrow(R A)+\operatorname{EXTS}\left(S I| |{ }^{16} 0\right)\)
The sum (RAIO) + (SIII 0x0000) is placed into register RT.

\section*{Special Registers Altered:}

None
Extended Mnemonics:
Examples of extended mnemonics for Add Immediate Shifted:
Extended:
lis \(\quad R x\), value
subis

Equivalent to: addis Rx,0,value addis \(R x, R y\), -value

\section*{Add PC Immediate Shifted DX-form}
addpcis RT,D
I
\begin{tabular}{|l|l|l|l|l|l|}
\hline 19 & RT & d 1 & d0 & 2 & 10 \\
\hline
\end{tabular}
\(\mathrm{D} \leftarrow \mathrm{d} 0||\mathrm{~d} 1|| \mathrm{d} 2\)
\(R T \leftarrow N I A+\operatorname{EXTS}\left(D| |{ }^{16} 0\right)\)
The sum of NIA + (D || \(0 \times 0000\) ) is placed into register RT.

Special Registers Altered:
None
Extended Mnemonics:
Examples of extended mnemonics for Add PC Immediate Shifted:
\begin{tabular}{lll} 
Extended: & \multicolumn{1}{l}{ Equivalent to: } \\
subpcis & \(R x\), value & addpcis \\
\(R x\), -value \\
Inia & \(R x\) & addpcis
\end{tabular}
XO-form
\[
\begin{aligned}
& (O E=0 \mathrm{Rc}=0) \\
& (\mathrm{OE}=0 \mathrm{Rc}=1) \\
& (\mathrm{OE}=1 \mathrm{Rc}=0) \\
& (\mathrm{OE}=1 \mathrm{Rc}=1)
\end{aligned}
\]
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{1}{|c|}{ RT } & RA & RB & OE & 266 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22
\end{tabular}
\(R T \leftarrow(R A)+(R B)\)
The sum (RA) + (RB) is placed into register RT.
Special Registers Altered:

CR0 SO OV OV32

\section*{I}

\section*{Add Immediate Carrying}

D-form
addic RT,RA,SI
\begin{tabular}{|l|l|l|lll|}
\hline 12 & RT & RA & \multicolumn{2}{|c|}{ SI } & 31 \\
\hline 0 & & 6 & 11 & 16 & \\
\hline
\end{tabular}
```

RT}\leftarrow(RA)+ EXTS(SI

```

The sum (RA) + SI is placed into register RT.
Special Registers Altered:

\section*{Extended Mnemonics:}

Example of extended mnemonics for Add Immediate Carrying:

\section*{Extended:}
subic \(R x, R y\),value

Equivalent to: addic \(R x, R y\),-value

\section*{Subtract From}

XO-form
( \(\mathrm{OE}=0 \mathrm{Rc}=0\) )
( \(\mathrm{OE}=0 \mathrm{Rc}=1\) )
( \(\mathrm{OE}=1 \mathrm{Rc}=0\) )
( \(\mathrm{OE}=1 \mathrm{Rc}=1\) )
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & OE & 40 & Rc \\
\hline 0 & 6 & & 11 & 16 & 21 & 22
\end{tabular}
\(R T \leftarrow\urcorner(R A)+(R B)+1\)
The sum \(\neg(R A)+(R B)+1\) is placed into register RT.
Special Registers Altered:
CRO OV OV32
SO OV
(if \(\mathrm{Rc}=1\) )
(if \(\mathrm{OE}=1\) )

\section*{Extended Mnemonics:}

Example of extended mnemonics for Subtract From:
\begin{tabular}{ll} 
Extended: & \multicolumn{1}{l}{ Equivalent to: } \\
sub \(\quad R x, R y, R z\) & subf \(R x, R z, R y\)
\end{tabular}

Add Immediate Carrying and Record
D-form
addic. RT,RA,SI
\begin{tabular}{|l|l|l|lll|}
\hline 13 & RT & RA & & SI & 31 \\
\hline
\end{tabular}
\(R T \leftarrow(R A)+\operatorname{EXTS}(S I)\)
The sum (RA) + SI is placed into register RT.
Special Registers Altered:
CRO CA CA32

\section*{Extended Mnemonics:}

Example of extended mnemonics for Add Immediate Carrying and Record:

Extended: subic. Rx,Ry,value

Equivalent to: addic. Rx,Ry,-value

\section*{Subtract From Immediate Carrying}

D-form
subfic RT,RA,SI
\begin{tabular}{|l|l|l|l|ll|}
\hline 8 & RT & RA & & SI & 31 \\
\hline 0 & & 6 & 11 & 16 & \\
\hline
\end{tabular}

RT \(\leftarrow \neg(\mathrm{RA})+\) EXTS \((\mathrm{SI})+1\)
The sum \(\neg(R A)+S I+1\) is placed into register RT.
Special Registers Altered:
I
CA CA32

\section*{Add Carrying \\ \begin{tabular}{ll} 
addc & RT,RA,RB \\
addc. & \(R T, R A, R B\) \\
addco & \(R T, R A, R B\) \\
addco. & RT,RA,RB
\end{tabular}}

\section*{XO-form}
\((O E=0 R c=0)\)
\((O E=0 R c=1)\)
\((O E=1 R c=0)\)
\((O E=1 R c=1)\)
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 31 & \multicolumn{1}{|c|}{ RT } & RA & RB & OE & & 10 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22 & \\
\hline
\end{tabular}
\(R T \leftarrow(R A)+(R B)\)
The sum (RA) + (RB) is placed into register RT.
Special Registers Altered:
CA CA32
CRO
SO OV OV32
\begin{tabular}{c|} 
(if \(\mathrm{Rc}=1\) ) \\
(if \(\mathrm{OE}=1\) )
\end{tabular}

Subtract From Carrying

\section*{XO-form}
\begin{tabular}{lll} 
subfc & \(R T, R A, R B\) & \((O E=0 R c=0)\) \\
subfc. & \(R T, R A, R B\) & \((O E=0 R c=1)\) \\
subfco & \(R T, R A, R B\) & \((O E=1 R c=0)\) \\
subfco. & \(R T, R A, R B\) & \((O E=1 R c=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|l|l|l|}
\hline 31 & \multicolumn{1}{c}{ RT } & RA & RB & OE & 8 & 8 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22 & & 31 \\
\hline
\end{tabular}
```

$R T \leftarrow \neg(R A)+(R B)+1$

```

The sum \(\neg(R A)+(R B)+1\) is placed into register RT.
Special Registers Altered:
\begin{tabular}{lr} 
I CA CA32 & \\
CRO OV OV32 & (if Rc=1) \\
(if \(\mathrm{OE}=1\) )
\end{tabular}

Example of extended mnemonics for Subtract From Carrying:
\begin{tabular}{ll} 
Extended: & \multicolumn{2}{l}{ Equivalent to: } \\
subc \(\quad R x, R y, R z\) & subfc \(\quad R x, R z, R y\)
\end{tabular}

Extended:
subc Rx,Ry,Rz
subfc Rx,Rz,Ry

\section*{Add Extended}
\begin{tabular}{lll} 
adde & RT,RA,RB & \((O E=0 \mathrm{Rc}=0)\) \\
adde. & RT,RA,RB & \((O E=0 \mathrm{Rc}=1)\) \\
addeo & RT,RA,RB & \((\mathrm{OE}=1 \mathrm{Rc}=0)\) \\
addeo. & RT,RA,RB & \((\mathrm{OE}=1 \mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{1}{|c|}{ RT } & RA & RB & OE & 138 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22
\end{tabular}
```

RT\leftarrow(RA) + (RB) + CA

```

The sum (RA) + (RB) + CA is placed into register RT.

\section*{Special Registers Altered:}

I
I SO OV OV32

\section*{Add to Minus One Extended}

XO-form
\begin{tabular}{lll} 
addme & RT,RA & \((O E=0 \mathrm{Rc}=0)\) \\
addme. & RT,RA & \((\mathrm{OE}=0 \mathrm{Rc}=1)\) \\
addmeo & RT,RA & \((\mathrm{OE}=1 \mathrm{Rc}=0)\) \\
addmeo. & RT,RA & \((\mathrm{OE}=1 \mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & \multicolumn{1}{|c|}{ III } & OE & 234 & Rc \\
\hline 0 & & 6 & 11 & 16 & & 22
\end{tabular}
\[
R T \leftarrow(R A)+C A-1
\]

The sum (RA) \(+C A+{ }^{64} 1\) is placed into register RT.
Special Registers Altered:

XO-form
\begin{tabular}{l|}
\(\underset{(\text { if } R c=1)}{ }\) \\
(if \(\mathrm{OE}=1)\)
\end{tabular}\(\quad\) |

Subtract From Extended
XO-form
\begin{tabular}{lll} 
subfe & \(R T, R A, R B\) & \((O E=0 \mathrm{Rc}=0)\) \\
subfe. & \(\mathrm{RT}, \mathrm{RA}, \mathrm{RB}\) & \((\mathrm{OE}=0 \mathrm{Rc}=1)\) \\
subfeo & \(\mathrm{RT}, \mathrm{RA}, \mathrm{RB}\) & \((\mathrm{OE}=1 \mathrm{Rc}=0)\) \\
subfeo. & \(\mathrm{RT}, \mathrm{RA}, \mathrm{RB}\) & \((\mathrm{OE}=1 \mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & OE & 136 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22 \\
\hline
\end{tabular}
\(R T \leftarrow \neg(R A)+(R B)+C A\)
The sum \(\neg(R A)+(R B)+C A\) is placed into register \(R T\).

\section*{Special Registers Altered:}

CA CA32
CRO
SO OV OV32
(if \(R c=1\) )
(if \(\mathrm{OE}=1\) )

Subtract From Minus One Extended XO-form
\begin{tabular}{lll} 
subfme & RT,RA & \((O E=0 \mathrm{Rc}=0)\) \\
subfme. & \(\mathrm{RT}, \mathrm{RA}\) & \((\mathrm{OE}=0 \mathrm{Rc}=1)\) \\
subfmeo & \(\mathrm{RT}, \mathrm{RA}\) & \((\mathrm{OE}=1 \mathrm{Rc}=0)\) \\
subfmeo. & \(\mathrm{RT}, \mathrm{RA}\) & \((\mathrm{OE}=1 \mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{1}{|c|}{ RT } & RA & \multicolumn{1}{|l|}{ I/I } & OE & 232 & Rc \\
0 & & & 11 & 16 & 21 & 22 \\
\hline
\end{tabular}
\(R T \leftarrow \neg(R A)+C A-1\)
The sum \(\neg(R A)+C A+{ }^{64} 1\) is placed into register RT.
Special Registers Altered:
CA CA32
CRO OV OV32
SO OV
(if \(\mathrm{Rc}=1\) )
(if \(\mathrm{OE}=1\) )

\section*{Add to Zero Extended}
\begin{tabular}{lll} 
addze & RT,RA & \((O E=0 \mathrm{Rc}=0)\) \\
addze. & RT,RA & \((\mathrm{OE}=0 \mathrm{Rc}=1)\) \\
addzeo & RT,RA & \((\mathrm{OE}=1 \mathrm{Rc}=0)\) \\
addzeo. & RT,RA & \((\mathrm{OE}=1 \mathrm{Rc}=1)\)
\end{tabular}

\section*{Subtract From Zero Extended}

\section*{XO-form}
\begin{tabular}{lll} 
subfze & RT,RA & \((O E=0 \quad R c=0)\) \\
subfze. & RT,RA & \((O E=0 \quad R c=1)\) \\
subfzeo & \(R T, R A\) & \((O E=1 R c=0)\) \\
subfzeo. & \(R T, R A\) & \((O E=1 R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & \multicolumn{1}{|c|}{ I/I } & OE & 200 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22 \\
31 \\
\hline
\end{tabular}
\(R T \leftarrow \neg(R A)+C A\)
The sum \(\neg(R A)+C A\) is placed into register RT.

\section*{Special Registers Altered:}

I
(if \(\mathrm{Rc}=1\) )

\section*{CA CA32}

CRO (if \(\mathrm{Rc}=1\) )
SO OV OV32
(if \(\mathrm{OE}=1\) )

\section*{Programming Note}

The setting of CA and CA32 by the Add and Subtract From instructions, including the Extended versions thereof, is mode-dependent. If a sequence of these instructions is used to perform extended-precision addition or subtraction, the same mode should be used throughout the sequence.

\section*{Negate}

XO-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline neg & \multicolumn{3}{|c|}{RT,RA} & \multicolumn{3}{|r|}{( \(\mathrm{OE}=0 \mathrm{Rc}=0\) )} \\
\hline neg. & \multicolumn{3}{|c|}{RT,RA} & \multicolumn{3}{|r|}{( \(\mathrm{OE}=0 \mathrm{Rc}=1\) )} \\
\hline nego & \multicolumn{3}{|c|}{RT,RA} & \multicolumn{3}{|r|}{( \(\mathrm{OE}=1 \mathrm{Rc}=0\) )} \\
\hline nego. & \multicolumn{3}{|c|}{RT,RA} & \multicolumn{3}{|r|}{( \(\mathrm{OE}=1 \mathrm{Rc}=1\) )} \\
\hline \[
31
\] & RT & \[
\int_{11} \mathrm{RA}
\] & \[
{ }_{16} / / I
\] & & 104 & Rc \\
\hline
\end{tabular}
\(R T \leftarrow \neg(R A)+1\)
The sum \(\neg(R A)+1\) is placed into register \(R T\).
If the processor is in 64-bit mode and register RA contains the most negative 64-bit number ( \(0 \times 8000\) _ 0000_0000_0000), the result is the most negative num-
\| ber and, if \(\mathrm{OE}=1, \mathrm{OV}\) and OV32 are set to 1 . Similarly, if the processor is in 32-bit mode and (RA) \({ }_{32: 63}\) contain the most negative 32-bit number ( \(0 \times 8000 \_0000\) ), the low-order 32 bits of the result contain the most negative
| 32-bit number and, if \(\mathrm{OE}=1, \mathrm{OV}\) and OV32 are set to 1 .

\section*{Special Registers Altered:}
\begin{tabular}{lr} 
CRO & (if \(\mathrm{Rc}=1\) ) \\
SO OV OV32 & (if \(\mathrm{OE}=1\) )
\end{tabular}

Multiply Low Immediate
D-form
mulli RT,RA,SI
\begin{tabular}{|l|l|l|l|lll|}
\hline 7 & RT & RA & \multicolumn{2}{|c|}{ SI } & 31 \\
\hline 0 & 6 & & 11 & 16 & & 3 \\
\hline
\end{tabular}
\(\operatorname{prod}_{0: 127} \leftarrow(\mathrm{RA}) \times \operatorname{EXTS}(S I)\)
\(\mathrm{RT} \leftarrow \operatorname{prod}_{64: 127}\)
The 64-bit first operand is (RA). The 64-bit second operand is the sign-extended value of the SI field. The low-order 64 bits of the 128-bit product of the operands are placed into register RT.

Both operands and the product are interpreted as signed integers.
```

Special Registers Altered:
None

```

\section*{Multiply Low Word}
```

XO-form

| mullw | RT,RA,RB | $(O E=0 \mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| mullw. | RT,RA,RB | $(\mathrm{OE}=0 \mathrm{Rc}=1)$ |
| mullwo | $\mathrm{RT}, \mathrm{RA}, \mathrm{RB}$ | $(\mathrm{OE}=1 \mathrm{Rc}=0)$ |
| mullwo. | RT,RA,RB | $(\mathrm{OE}=1 \mathrm{Rc}=1)$ |

```
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & OE & 235 & Rc \\
0 & & & 11 & 16 & 21 & 22 \\
\hline
\end{tabular}
```

RT}\leftarrow(RA) 32:63 ` (RB) 32:63

```

The 32-bit operands are the low-order 32 bits of RA and of RB. The 64-bit product of the operands is placed into register RT.
If \(\mathrm{OE}=1\) then OV and OV32 are set to 1 if the product cannot be represented in 32 bits.

Both operands and the product are interpreted as signed integers.
Special Registers Altered:
I \begin{tabular}{lr} 
CRO & (if \(\mathrm{Rc}=1\) ) \\
SO OV OV32 & (if \(\mathrm{OE}=1\) )
\end{tabular}

\section*{- Programming Note}

For mulli and mullw, the low-order 32 bits of the product are the correct 32 -bit product for 32 -bit mode.

For mulli and mulld, the low-order 64 bits of the product are independent of whether the operands are regarded as signed or unsigned 64-bit integers. For mulli and mullw, the low-order 32 bits of the product are independent of whether the operands are regarded as signed or unsigned 32-bit integers.

\section*{Multiply High Word}

XO-form
\begin{tabular}{lll} 
mulhw & RT,RA,RB & \((\mathrm{Rc}=0)\) \\
mulhw. & \(\mathrm{RT}, \mathrm{RA}, \mathrm{RB}\) & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & \(/\) & & 75 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22 & \\
31 \\
\hline
\end{tabular}
\(\operatorname{prod}_{0: 63} \leftarrow(\mathrm{RA})_{32: 63} \times(\mathrm{RB})_{32: 63}\)
\(\mathrm{RT}_{32: 63} \leftarrow \operatorname{prod}_{0: 31}\)
\(\mathrm{RT}_{0: 31} \leftarrow\) undefined
The 32-bit operands are the low-order 32 bits of RA and of RB. The high-order 32 bits of the 64-bit product of the operands are placed into \(\mathrm{RT}_{32: 63}\). The contents of \(R T_{0: 31}\) are undefined.
Both operands and the product are interpreted as signed integers.

\section*{Special Registers Altered:}

CRO (bits 0:2 undefined in 64-bit mode) (if \(\mathrm{Rc}=1\) )

\section*{Multiply High Word Unsigned}

XO-form
\begin{tabular}{lll} 
mulhwu & RT,RA,RB & \((\mathrm{Rc}=0)\) \\
mulhwu. & \(\mathrm{RT}, \mathrm{RA}, \mathrm{RB}\) & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & 1 & & 11 \\
0 & & & 11 & 16 & 21 & 22 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \operatorname{prod}_{0: 63} \leftarrow(\mathrm{RA})_{32: 63} \times(\mathrm{RB})_{32: 63} \\
& \mathrm{RT}_{32: 63} \leftarrow \text { prod }_{0: 31} \\
& \mathrm{RT}_{0: 31} \leftarrow \text { undefined }
\end{aligned}
\]

The 32-bit operands are the low-order 32 bits of RA and of RB. The high-order 32 bits of the 64-bit product of the operands are placed into \(\mathrm{RT}_{32: 63}\). The contents of \(R T_{0: 31}\) are undefined.
Both operands and the product are interpreted as unsigned integers, except that if Rc=1 the first three bits of CR Field 0 are set by signed comparison of the result to zero.

\section*{Special Registers Altered:}

CRO (bits 0:2 undefined in 64-bit mode) (if Rc=1)

\section*{Divide Word}
\begin{tabular}{lll} 
divw & RT,RA,RB & \((O E=0 \mathrm{Rc}=0)\) \\
divw. & RT,RA,RB & \((O E=0 \mathrm{Rc}=1)\) \\
divwo & RT,RA,RB & \((O E=1 \mathrm{Rc}=0)\) \\
divwo. & RT,RA,RB & \((O E=1 \mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{1}{|c|}{ RT } & RA & RB & OE & 491 & Rc \\
0 & & & 11 & 16 & 21 & 22 \\
31 \\
\hline
\end{tabular}
dividend \(_{0: 31} \leftarrow(\text { RA })_{32: 63}\)
divisor \(_{0: 31} \leftarrow(\mathrm{RB})_{32: 63}\)
\(\mathrm{RT}_{32: 63} \leftarrow\) dividend \(\div\) divisor
\(\mathrm{RT}_{0: 31} \leftarrow\) undefined
The 32-bit dividend is \((R A)_{32: 63}\). The 32-bit divisor is \((\mathrm{RB})_{32: 63}\). The 32-bit quotient is placed into \(\mathrm{RT}_{32: 63}\). The contents of \(R T_{0: 31}\) are undefined. The remainder is not supplied as a result.

Both operands and the quotient are interpreted as signed integers. The quotient is the unique signed integer that satisfies
\[
\text { dividend }=(\text { quotient } \times \text { divisor })+r
\]
where \(0 \leq r<\) Idivisorl if the dividend is nonnegative, and -|divisorl \(<r \leq 0\) if the dividend is negative.

If an attempt is made to perform any of the divisions
```

0x8000_0000 \div-1
<anything> \div 0

```
then the contents of register RT are undefined as are (if \(R c=1\) ) the contents of the LT, GT, and EQ bits of \(C R\)
| Field 0 . In these cases, if \(\mathrm{OE}=1\) then OV and OV32 are set to 1 .

Special Registers Altered:
CR0 (bits 0:2 undefined in 64-bit mode) (if Rc=1)
SO OV OV32
(if \(\mathrm{OE}=1\) )

\section*{Programming Note}

The 32-bit signed remainder of dividing (RA) 32:63 \(^{2}\) by (RB) 32:63 can be computed as follows, except in the case that \((R A)_{32: 63}=-2^{31}\) and \((R B)_{32: 63}=-1\).
\[
\begin{array}{lll}
\text { divw } & \text { RT, RA, RB } & \text { \# RT }=\text { quotient } \\
\text { mullw } & \text { RT, RT, RB } & \text { \# RT }=\text { quotient } \times \text { divisor } \\
\text { subf } & \text { RT, RT, RA } & \text { \# RT }=\text { remainder }
\end{array}
\]

Divide Word Unsigned
\begin{tabular}{lll} 
divwu & RT,RA,RB & \((O E=0 \mathrm{Rc}=0)\) \\
divwu. & RT,RA,RB & \((O E=0 \mathrm{Rc}=1)\) \\
divwuo & RT,RA,RB & \((O E=1 \mathrm{Rc}=0)\) \\
divwuo. & RT,RA,RB & \((O E=1 \mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & OE & 459 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22
\end{tabular}
dividend \(_{0: 31} \leftarrow(R A)_{32: 63}\)
divisor \(0: 31 \leftarrow(\mathrm{RB})_{32: 63}\)
\(\mathrm{RT}_{32: 63} \leftarrow\) dividend \(\div\) divisor
\(\mathrm{RT}_{0: 31} \leftarrow\) undefined
The 32 bit dividend is \((R A)_{32: 63}\). The 32-bit divisor is \((\mathrm{RB})_{32: 63}\). The 32-bit quotient is placed into \(\mathrm{RT}_{32: 63}\). The contents of \(R T_{0: 31}\) are undefined. The remainder is not supplied as a result.

Both operands and the quotient are interpreted as unsigned integers, except that if Rc=1 the first three bits of CR Field 0 are set by signed comparison of the result to zero. The quotient is the unique unsigned integer that satisfies
\[
\text { dividend }=(\text { quotient } \times \text { divisor })+r
\]
where \(0 \leq r<\) divisor.
If an attempt is made to perform the division
\[
\text { <anything> } \div 0
\]
then the contents of register RT are undefined as are (if \(R \mathrm{c}=1\) ) the contents of the LT, GT, and EQ bits of CR Field 0 . In this case, if \(\mathrm{OE}=1\) then OV and OV 32 are set to 1.

\section*{Special Registers Altered:}
\[
\text { CRO (bits 0:2 undefined in 64-bit mode) (if } R c=1 \text { ) }
\] SO OV OV32
(if \(\mathrm{OE}=1\) )

\section*{Programming Note}

The 32-bit unsigned remainder of dividing (RA) \({ }_{32: 63}\) by \((\mathrm{RB})_{32: 63}\) can be computed as follows.
\begin{tabular}{lll} 
divwu & RT, RA, RB & \# RT \(=\) quotient \\
mullw & RT, RT, RB & \# RT \(=\) quotient×divisor \\
subf & RT, RT, RA & \# RT \(=\) remainder
\end{tabular}

\section*{Modulo Signed Word X-form}
modsw
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT,RA,RB \\
0 & & & RT & RA & RB & \\
\hline 11 & & 779 & \\
\hline 31 \\
\hline
\end{tabular}
```

dividend }\leftarrow\mathrm{ EXTS((RA) 32:63)
divisor }\leftarrow\mathrm{ EXTS((RB)32:63)
remainder }\leftarrow\mathrm{ dividend % divisor
RT \leftarrowChop(remainder, 64)

```

Let di vidend be the signed integer word in bits 32:63 of register RA.

Let divisor be the signed integer word in bits 32:63 of register \(R B\).

The remainder of dividend divided by divisor is placed into register RT. The quotient is not supplied as a result.

The remainder is the unique signed integer that satisfies
remainder = dividend \(\cdot\) (quotient \(x\) divisor)
where \(0 \leq r e m a i n d e r<\mid\) divisor| if the dividend is nonnegative, and \(\cdot \mid\) divisor \(\mid<r e m a i n d e r \leq 0\) if the dividend is negative.

If an attempt is made to perform any of the divisions
\[
\begin{gathered}
0 \times 80000000 \% \cdot 1 \\
\text { <anvthina> \% } 0 \text { ol }
\end{gathered}
\]
then the contents of register RT are undefined.

\section*{Special Registers Altered: \\ None}

\section*{Modulo Unsigned Word X-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{moduw RT,RA,RB} \\
\hline \[
{ }_{0} 31
\] & \[
{ }_{6} \mathrm{RT}
\] & \[
{ }_{11} \mathrm{RA}
\] & \[
{ }_{16} \mathrm{RB}
\] & \[
\begin{array}{ll} 
& 267 \\
21 &
\end{array}
\] & \(|\)\begin{tabular}{|r|}
1 \\
31
\end{tabular} \\
\hline
\end{tabular}
```

dividend }\leftarrow\mathrm{ EXTZ((RA) 32:63)
divisor }\leftarrow\mathrm{ EXTZ((RB) 32:63)
remainder }\leftarrow\mathrm{ dividend % divisor
RT }\leftarrow\mathrm{ Chop(remainder, 64)

```

Let di vidend be the unsigned integer word in bits 32:63 of register RA.

Let divisor be the unsigned integer word in bits 32:63 of register RB.

The remainder of dividend divided by divisor is placed into register RT. The quotient is not supplied as a result.

The remainder is the unique unsigned integer that satisfies
\[
\text { remainder }=\text { dividend } \cdot \text { (quotient } \times \text { divisor) }
\]
where 0 sremainder < divisor.
If an attempt is made to perform any of the divisions
<anything> \% o
then the contents of register RT are undefined.

\section*{Special Registers Altered:}

None

\section*{Divide Word Extended}
\begin{tabular}{lll} 
divwe & RT,RA,RB & \((O E=0 \mathrm{Rc}=0)\) \\
divwe. & RT,RA,RB & \((O E=0 \mathrm{Rc}=1)\) \\
divweo & RT,RA,RB & \((O E=1 \mathrm{Rc}=0)\) \\
divweo. & RT,RA,RB & \((O E=1 \mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & \({ }^{2} \mathrm{RT}\) & \({ }_{11} \mathrm{RA}\) & \({ }_{16} \mathrm{RB}\) & OE & 427 & Rc \\
0 & & 21 & 22 & 31 \\
\hline
\end{tabular}
dividend \(_{0: 63} \leftarrow(\mathrm{RA})_{32: 63}| |^{32} 0\)
divisor \(_{0: 31} \leftarrow(\mathrm{RB})_{32: 63}\)
\(\mathrm{RT}_{32: 63} \leftarrow\) dividend \(\div\) divisor
\(\mathrm{RT}_{0: 31} \leftarrow\) undefined
The 64-bit dividend is (RA) 32:63 \({ }^{11}{ }^{32} 0\). The 32-bit divisor is \((R B)_{32: 63}\). If the quotient can be represented in 32 bits, it is placed into \(\mathrm{RT}_{32: 63}\). The contents of \(\mathrm{RT}_{0: 31}\) are undefined. The remainder is not supplied as a result.

Both operands and the quotient are interpreted as signed integers. The quotient is the unique signed integer that satisfies
\[
\text { dividend }=(\text { quotient } \times \text { divisor })+r
\]
where \(0 \leq r<\) Idivisorl if the dividend is nonnegative, and -|divisorl \(<r \leq 0\) if the dividend is negative.

If the quotient cannot be represented in 32 bits, or if an attempt is made to perform the division
```

<anything> \div 0

```
then the contents of register RT are undefined as are (if \(\mathrm{Rc}=1\) ) the contents of the LT, GT, and EQ bits of CR I Field 0 . In these cases, if \(\mathrm{OE}=1\) then OV and OV 32 are set to 1.

\section*{Special Registers Altered:}

CRO (bits 0:2 undefined in 64-bit mode) (if \(\mathrm{Rc}=1\) )
SO OV OV32
(if \(\mathrm{OE}=1\) )

XO-form
divweu
divweu.
divweuo divweuo.

RT,RA,RB
RT,RA,RB
RT,RA,RB
RT,RA,RB
( \(\mathrm{OE}=0 \mathrm{Rc}=1\) )
( \(\mathrm{OE}=1 \mathrm{Rc}=0\) )
(OE=1 Rc=1)
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & OE & 395 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22
\end{tabular}
dividend \(_{0: 63} \leftarrow(\mathrm{RA})_{32: 63}| |^{32} 0\)
divisor \(_{0: 31} \leftarrow(\mathrm{RB})_{32: 63}\)
\(\mathrm{RT}_{32: 63} \leftarrow\) dividend \(\div\) divisor
\(\mathrm{RT}_{0: 31} \leftarrow\) undefined
The 64-bit dividend is (RA) \()_{32: 63}{ }^{11}{ }^{32} 0\). The 32-bit divisor is \((R B)_{32: 63}\). If the quotient can be represented in 32 bits, it is placed into \(\mathrm{RT}_{32: 63}\). The contents of \(\mathrm{RT}_{0: 31}\) are undefined. The remainder is not supplied as a result.

Both operands and the quotient are interpreted as unsigned integers, except that if Rc=1 the first three bits of CR Field 0 are set by signed comparison of the result to zero. The quotient is the unique unsigned integer that satisfies
\[
\text { dividend }=(\text { quotient } \times \text { divisor })+r
\]
where \(0 \leq r<\) divisor.
If \((R A) \geq(R B)\), or if an attempt is made to perform the division
```

<anything> \div 0

```
then the contents of register RT are undefined as are (if \(\mathrm{Rc}=1\) ) the contents of the LT, GT, and EQ bits of CR - Field 0 . In these cases, if \(\mathrm{OE}=1\) then OV and OV 32 are set to 1 .

\section*{Special Registers Altered:}

CRO (bits 0:2 undefined in 64-bit mode) (if \(\mathrm{Rc}=1\) )
SO OV OV32 (if \(\mathrm{OE}=1\) )

\section*{Version 3.0}

\section*{Programming Note}

Unsigned long division of a 64-bit dividend contained in two 32-bit registers by a 32-bit divisor can be computed as follows. The algorithm is shown first, followed by Assembler code that implements the algorithm. The dividend is Dh II DI , the divisor is Dv , and the quotient and remainder are \(Q\) and \(R\) respectively, where these variables and all intermediate variables represent unsigned 32-bit integers. It is assumed that Dv > Dh, and that assigning a value to an intermediate variable assigns the low-order 32 bits of the value and ignores any higher-order bits of the value. (In both the algorithm and the Assembler code, " r 1 " and " r 2 " refer to "remainder 1" and "remainder 2", rather than to GPRs 1 and 2.)

\section*{Algorithm:}
3. \(q 1 \leftarrow\) divweu Dh, Dv
4. \(\mathrm{r} 1 \leftarrow-(\mathrm{q} 1 \times \mathrm{Dv}) \quad\) \# remainder of step 1

> divide operation
> (see Note 1)
5. \(q 2 \leftarrow d i v w u \operatorname{DI}, \mathrm{Dv}\)
6. \(\mathrm{r} 2 \leftarrow \mathrm{DI}-(\mathrm{q} 2 \times \mathrm{Dv}) \quad\) \# remainder of step 2 divide operation
7. \(Q \leftarrow q 1+q 2\)
8. \(R \leftarrow r 1+r 2\)
9. if \((R<r 2) \mid(R \geq D v)\) then \(\quad\) (see Note 2)
\(Q \leftarrow Q+1\) \# increment quotient
\(R \leftarrow R-D v\) \# decrement rem'der

Assembler Code:
\begin{tabular}{|c|c|c|}
\hline divweu & r3, r4,r6 & \# q1 \\
\hline divwu & r7, r5,r6 & \# q2 \\
\hline mullw & r8, r3, r6 & \# -r1 = q1 * Dv \\
\hline mullw & r0, r7, r6 & \# q2 * Dv \\
\hline subf & r10,r0, r5 & \# r2 = Dl - (q2 * Dv) \\
\hline add & r3, r3, r7 & \# Q = q1 + q2 \\
\hline subf & r4,r8,r10 & \# \(\mathrm{R}=\mathrm{r} 1+\mathrm{r} 2\) \\
\hline cmplw & r4,r10 & \# \(\mathrm{R}<\mathrm{r} 2\) ? \\
\hline blt & *+12 & \# must adjust Q and R if yes \\
\hline cmplw & r4,r6 & \# \(\mathrm{R} \geq \mathrm{Dv}\) ? \\
\hline blt & *+12 & \# must adjust Q and R if yes \\
\hline addi & r3, r3,1 & \# Q = Q + 1 \\
\hline subf & r4,r6,r4 & \# R = R - Dv \\
\hline \multicolumn{3}{|l|}{\multirow[t]{2}{*}{\begin{tabular}{l}
\# Quotient in r3 \\
\# Remainder in r4
\end{tabular}}} \\
\hline & & \\
\hline
\end{tabular}

\section*{Notes:}
1. The remainder is \(\mathrm{Dh}\left\|\|^{32} 0-(q 1 \times \mathrm{Dv})\right.\). Because the remainder must be less than Dv and \(D v<2^{32}\), the remainder is representable in 32 bits. Because the low-order 32 bits of \(\mathrm{Dh} \|{ }^{32} 0\) are 0 s, the remainder is therefore equal to the low-order 32 bits of \(-(q 1 \times\) Dv). Thus assigning \(-(q 1 \times D v)\) to \(r 1\) yields the correct remainder.
2. \(R\) is less than \(r 2\) (and also less than \(r 1\) ) if and only if the addition at step 6 carried out of 32 bits - i.e., if and only if the correct sum could not be represented in 32 bits - in which case the correct sum is necessarily greater than Dv.
3. For additional information see the book Hacker's Delight, by Henry S. Warren, Jr., as potentially amended at the web site http://www.hackersdelight.org.

darn RT,L
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline 31 & & RT & & I/I & L & & I/I & \\
\hline 0 & & 6 & & 11 & 13 & 14 & 16 & \\
\hline
\end{tabular}
\(\mathrm{RT} \leftarrow \operatorname{random}(\mathrm{L})\)
A random number is placed into register RT in a format selected by L as shown in the following table. The value 0xFFFFFFFFF_FFFFFFFF indicates an error condition. For \(\mathrm{L}=0\), the random number range is \(0: 0 x F F F F F F F F\). For \(L=1\) and \(L=2\), the random number range is \(0: 0 x F F F F F F F F\) _FFFFFFFFE.
\begin{tabular}{|l|l|}
\hline \(\mathbf{L}\) & Format \\
\hline 0 & \({ }^{32} 0\) II \(\mathrm{CRN}_{0: 31}\) \\
\hline 1 & \(\mathrm{CRN}_{0: 63}\) \\
\hline 2 & \(\mathrm{RRN}_{0: 63}\) \\
\hline 3 & reserved \\
\hline
\end{tabular}

Format above is for non-error conditions.
0xFFFFFFFFF_FFFFFFFF for error conditions.
CRN \(=\) conditioned random number
RRN = raw random number

A raw random number is unconditioned noise source output. A conditioned random number has been processed by hardware to reduce bias.

\section*{Special Registers Altered:}
none

\section*{Programming Note}

32-bit software running in an environment that does not preserve the high-order 32 bits of GPRs across invocations of the system error handler, signal handlers, event-based branch handlers, etc. may use the \(\mathrm{L}=0\) variant of darn and interpret the value \(0 x F F F F F F F F\) to indicate an error condition. The fact that the error condition includes the valid value 0x00000000_FFFFFFFFF together with the true error value 0xFFFFFFFFF_FFFFFFFF is not a problem.

\section*{Programming Note}

When the error value is obtained, software is expected to repeat the operation. If a non-error value has not been obtained after several attempts, a software random number generation method should be used. The recommended number of attempts may be implementation specific. In the absence of other guidance, ten attempts should be adequate.

\section*{Programming Note}

The random number generator provided by this instruction is NIST SP800-90B and SP800-90C compliant to the extent possible given the completeness of the standards at the time the hardware is designed. The random number generator provides a minimum of 0.5 bits of entropy per bit.

\subsection*{3.3.9.1 64-bit Fixed-Point Arithmetic Instructions}

\section*{Multiply Low Doubleword}
\begin{tabular}{lll} 
mulld & RT,RA,RB & \((O E=0 \mathrm{Rc}=0)\) \\
mulld. & RT,RA,RB & \((O E=0 \mathrm{Rc}=1)\) \\
mulldo & RT,RA,RB & \((O E=1 \mathrm{Rc}=0)\) \\
mulldo. & RT,RA,RB & \((O E=1 \mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 31 & RT & RA & RB & OE & 233 & Rc \\
0 & & & 11 & 16 & 21 & 22 & 31 \\
\hline
\end{tabular}
```

prod}0:127 \leftarrow(RA) × (RB
RT}\leftarrow\mp@subsup{\operatorname{prod}}{64:127}{

```

The 64-bit operands are (RA) and (RB). The low-order 64 bits of the 128 -bit product of the operands are placed into register RT.
| If \(\mathrm{OE}=1\) then OV and OV 32 are set to 1 if the product cannot be represented in 64 bits.

Both operands and the product are interpreted as signed integers.

\section*{Special Registers Altered:}

\section*{CR0 \\ SO OV OV32}
(if \(\mathrm{Rc}=1\) )
(if \(\mathrm{OE}=1\) )

\section*{Programming Note}

The XO-form Multiply instructions may execute faster on some implementations if RB contains the operand having the smaller absolute value.

Multiply High Doubleword
XO-form

\(\operatorname{prod}_{0: 127} \leftarrow(\mathrm{RA}) \times(\mathrm{RB})\)
\(R T \leftarrow\) prod \(_{0: 63}\)
The 64-bit operands are (RA) and (RB). The high-order 64 bits of the 128 -bit product of the operands are placed into register RT.

Both operands and the product are interpreted as signed integers.

\section*{Special Registers Altered:}

CRO
(if \(\mathrm{Rc}=1\) )

\section*{Multiply High Doubleword Unsigned}

XO-form
mulhdu mulhdu.

( \(\mathrm{Rc}=0\) ) ( \(\mathrm{Rc}=1\) )
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & \(/\) & & 9 \\
\hline 0 & & 6 & 11 & 16 & 21 & 22 \\
\hline
\end{tabular}
\(\operatorname{prod}_{0: 127} \leftarrow(\mathrm{RA}) \times(\mathrm{RB})\)
RT \(\leftarrow \operatorname{prod}_{0: 63}\)
The 64-bit operands are (RA) and (RB). The high-order 64 bits of the 128 -bit product of the operands are placed into register RT.

Both operands and the product are interpreted as unsigned integers, except that if Rc=1 the first three bits of CR Field 0 are set by signed comparison of the result to zero.
Special Registers Altered:
CRO
(if \(\mathrm{Rc}=1\) )

\section*{Multiply-Add High Doubleword VA-form}
maddhd RT,RA.RB,RC

```

result \& (EXTS(GPR[RA]) x EXTS(GPR[RB])) + EXTS(GPR[RC])
GPR[RT].dword[0]}\leftarrowChop(result>>64,64

```

The signed integer value in GPR[RA] is multiplied by the signed integer value in GPR[RB]. The 128-bit product is added to the signed integer value in GPR[RC]. The upper 64 bits of the result are placed into GPR[ RT] .

\section*{Special Registers Altered:}

None

\section*{Multiply-Add High Doubleword Unsigned VA-form}
maddhdu RT,RA.RB,RC
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 4 & \[
\begin{array}{ll} 
& \mathrm{RT}
\end{array}
\] & \[
{ }_{11} \quad \mathrm{RA}
\] & \[
{ }_{16} \mathrm{RB}
\] & & RC & 26 & 31 \\
\hline \multicolumn{8}{|l|}{result \(\leftarrow(\) EXTZ(GPR[RA]) \(\times\) EXTZ(GPR[RB]) ) + EXTZ(GPR[RC])} \\
\hline \multicolumn{8}{|l|}{GPR[RT].dword[0] ¢Chop(result \(\ggg 64,64\) )} \\
\hline
\end{tabular}

The unsigned integer value in GPR[RA] is multiplied by the unsigned integer value in GPR[RB]. The 128 -bit product is added to the unsigned integer value in GPR[RC]. The upper 64 bits of the result are placed into GPR[RT].

Special Registers Altered:
None

\section*{Multiply-Add Low Doubleword VA-form}
maddld RT,RA.RB,RC
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & \[
\sigma_{6} \mathrm{RT}
\] & \({ }_{11} R A\) & \({ }_{16} \mathrm{RB}\) & \({ }_{21} \mathrm{RC}\) & 26 & 31 \\
\hline \multicolumn{7}{|l|}{} \\
\hline \multicolumn{7}{|l|}{GPR[RT].dwordi0] ¢ Chop(result, 64)} \\
\hline
\end{tabular}

The integer value in GPR[RA] is multiplied by the integer value in GPR[RB]. The 128 -bit product is added to the integer value in GPR[ RC]. The lower 64 bits of the result are placed into GPR[RT].

Special Registers Altered:
None

\section*{Divide Doubleword}

XO-form
\begin{tabular}{lll} 
divd & \(R T, R A, R B\) & \((O E=0 \quad R c=0)\) \\
divd. & \(R T, R A, R B\) & \((O E=0 \quad R c=1)\) \\
divdo & \(R T, R A, R B\) & \((O E=1 R c=0)\) \\
divdo. & \(R T, R A, R B\) & \((O E=1 R c=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|c|c|c|c|c|}
\hline 31 & RT & RA & RB & OE & 489 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22 & 31 \\
\hline
\end{tabular}
```

dividend 0:63}\leftarrow(RA
divisor0:63 }\leftarrow(RB
RT}\leftarrow\mathrm{ dividend }\div\mathrm{ divisor

```

The 64-bit dividend is (RA). The 64-bit divisor is (RB). The 64-bit quotient is placed into register RT. The remainder is not supplied as a result.

Both operands and the quotient are interpreted as signed integers. The quotient is the unique signed integer that satisfies
\[
\text { dividend }=(q u o t i e n t \times \text { divisor })+r
\]
where \(0 \leq r<\) Idivisorl if the dividend is nonnegative, and - Idivisorl < \(r \leq 0\) if the dividend is negative.

If an attempt is made to perform any of the divisions
```

0x8000_0000_0000_0000 \div-1
<anything> - % 0

```
then the contents of register RT are undefined as are (if Rc=1) the contents of the LT, GT, and EQ bits of CR Field 0. In these cases, if \(\mathrm{OE}=1\) then OV and OV 32 are set to 1 .

\section*{Special Registers Altered:}

\section*{CRO \\ SO OV OV32}
(if \(\mathrm{Rc}=1\) ) (if \(\mathrm{OE}=1\) )

\section*{Programming Note}

The 64-bit signed remainder of dividing (RA) by (RB) can be computed as follows, except in the case that \((R A)=-2^{63}\) and \((R B)=-1\).
\[
\begin{array}{lll}
\text { divd } & \text { RT, RA, RB } & \text { \# RT }=\text { quotient } \\
\text { mulld } & R T, R T, R B & \text { \# RT }=\text { quotient } \times \text { divisor } \\
\text { subf } & \text { RT, RT, RA } & \text { \# RT }=\text { remainder }
\end{array}
\]

\section*{Divide Doubleword Unsigned}

\author{
XO-form
}
\begin{tabular}{lll} 
divdu & \(R T, R A, R B\) & \((O E=0 \mathrm{Rc}=0)\) \\
divdu. & \(\mathrm{RT}, \mathrm{RA}, \mathrm{RB}\) & \((\mathrm{OE}=0 \mathrm{Rc}=1)\) \\
divduo & \(\mathrm{RT}, \mathrm{RA}, \mathrm{RB}\) & \((\mathrm{OE}=1 \mathrm{Rc}=0)\) \\
divduo. & \(\mathrm{RT}, \mathrm{RA}, \mathrm{RB}\) & \((\mathrm{OE}=1 \mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & OE & \multicolumn{1}{|c|}{457} & Rc \\
0 & & & 11 & 16 & 21 & 22
\end{tabular}
dividend \(_{0: 63} \leftarrow(\) RA \()\)
divisor \(0: 63 \leftarrow\) (RB)
\(\mathrm{RT} \leftarrow\) dividend \(\div\) divisor
The 64-bit dividend is (RA). The 64-bit divisor is (RB). The 64-bit quotient is placed into register RT. The remainder is not supplied as a result.

Both operands and the quotient are interpreted as unsigned integers, except that if Rc=1 the first three bits of CR Field 0 are set by signed comparison of the result to zero. The quotient is the unique unsigned integer that satisfies
\[
\text { dividend }=(\text { quotient } \times \text { divisor })+r
\]
where \(0 \leq r<\) divisor.
If an attempt is made to perform the division
```

<anything> \div 0

```
then the contents of register RT are undefined as are (if \(R c=1\) ) the contents of the LT, GT, and EQ bits of CR Field 0 . In this case, if \(\mathrm{OE}=1\) then OV and OV32 are set to 1.

Special Registers Altered:

CRO
SO OV OV32
(if \(\mathrm{Rc}=1\) ) (if \(\mathrm{OE}=1\) )

\section*{Programming Note}

The 64-bit unsigned remainder of dividing (RA) by (RB) can be computed as follows.
```

divdu RT,RA,RB \# RT = quotient
mulld RT,RT,RB \# RT = quotientxdivisor
subf RT,RT,RA \# RT = remainder

```

\section*{Divide Doubleword Extended}

XO-form
\begin{tabular}{lll} 
divde & \(R T, R A, R B\) & \((O E=0 R c=0)\) \\
divde. & \(R T, R A, R B\) & \((O E=0 \quad R c=1)\) \\
divdeo & \(R T, R A, R B\) & \((O E=1 R c=0)\) \\
divdeo. & \(R T, R A, R B\) & \((O E=1 R c=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & RT & RA & RB & OE & 425 & Rc \\
0 & & & 11 & 16 & 21 & 22 \\
\hline
\end{tabular}
dividend \({ }_{0: 127} \leftarrow(\mathrm{RA})| |{ }^{64} 0\)
divisor \(0: 63 \leftarrow\) (RB)
RT \(\leftarrow\) dividend \(\div\) divisor
The 128 -bit dividend is (RA) II \({ }^{64} 0\). The 64-bit divisor is (RB). If the quotient can be represented in 64 bits, it is placed into register RT. The remainder is not supplied as a result.

Both operands and the quotient are interpreted as signed integers. The quotient is the unique signed integer that satisfies
\[
\text { dividend }=(\text { quotient } \times \text { divisor })+r
\]
where \(0 \leq r<\) Idivisorl if the dividend is nonnegative, and -Idivisorl \(<r \leq 0\) if the dividend is negative.

If the quotient cannot be represented in 64 bits, or if an attempt is made to perform the division
```

<anything> $\div 0$

```
then the contents of register RT are undefined as are (if \(\mathrm{Rc}=1\) ) the contents of the LT, GT, and EQ bits of CR
| Field 0. In these cases, if \(\mathrm{OE}=1\) then OV and OV32 are set to 1.

\section*{Special Registers Altered:}

CRO
I
(if Rc=1)
(if \(\mathrm{OE}=1\) )

Divide Doubleword Extended Unsigned
XO-form
\begin{tabular}{lll} 
divdeu & RT,RA,RB & \((O E=0 \quad R c=0)\) \\
divdeu. & \(R T, R A, R B\) & \((O E=0 \quad R c=1)\) \\
divdeuo & \(R T, R A, R B\) & \((O E=1 R c=0)\) \\
divdeuo. & \(R T, R A, R B\) & \((O E=1 R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{1}{|c|}{ RT } & RA & RB & OE & 393 & Rc \\
0 & & 6 & 11 & 16 & 21 & 22
\end{tabular}
dividend \(0_{0: 127} \leftarrow(\) RA \() \|{ }^{64} 0\)
divisor \(_{0: 63} \leftarrow\) (RB)
\(\mathrm{RT} \leftarrow\) dividend \(\div\) divisor
The 128 -bit dividend is (RA) || \({ }^{64} 0\). The 64-bit divisor is (RB). If the quotient can be represented in 64 bits, it is placed into register RT. The remainder is not supplied as a result.

Both operands and the quotient are interpreted as unsigned integers, except that if Rc=1 the first three bits of CR Field 0 are set by signed comparison of the result to zero. The quotient is the unique unsigned integer that satisfies
\[
\text { dividend }=(\text { quotient } \times \text { divisor })+r
\]
where \(0 \leq r<\) divisor.
If \((R A) \geq(R B)\), or if an attempt is made to perform the division
\[
\text { <anything> } \div 0
\]
then the contents of register RT are undefined as are (if \(\mathrm{Rc}=1\) ) the contents of the LT, GT, and EQ bits of CR
| Field 0. In these cases, if \(\mathrm{OE}=1\) then OV and OV32 are set to 1 .

Special Registers Altered:
\begin{tabular}{lr} 
CRO & (if \(\mathrm{Rc}=1)\) \\
SO OV OV32 & (if \(\mathrm{OE}=1\) )
\end{tabular}

\section*{Programming Note}

Unsigned long division of a 128-bit dividend contained in two 64-bit registers by a 64-bit divisor can be accomplished using the technique described in the Programming Note with the divweu instruction description: divd[e]u would be used instead of divw[e]u (and cmpld instead of cmplw, etc.).

\section*{Modulo Signed Doubleword X-form}
modsd
RT,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & & 777 & \begin{tabular}{r}
\(\prime\) \\
31 \\
\hline
\end{tabular}
\end{tabular}
```

dividend }\leftarrow\mathrm{ EXTS((RA))
divisor }\leftarrow\mathrm{ EXTS((RB))
remainder }\leftarrow\mathrm{ dividend % divisor
RT \leftarrowChop(remainder, 64)

```

Let dividend be the signed integer doubleword in register RA.

Let divisor be the signed integer doubleword in register \(R B\).

The remainder of dividend divided by divisor is placed into register RT. The quotient is not supplied as a result.

The remainder is the unique signed integer that satisfies
```

remainder = dividend . (quotient x divisor)

```
where \(0 \leq r e m a i n d e r ~<~|d i v i s o r| ~ i f ~ t h e ~ d i v i d e n d ~ i s ~\) nonnegative, and • |divisor| < remainder \(\leq 0\) if the dividend is negative.

If an attempt is made to perform any of the divisions
```

<anything> % 0

```
then the contents of register RT are undefined.

\section*{Special Registers Altered:}

None

Modulo Unsigned Doubleword X-form
modud RT,RA,RB

```

dividend }\leftarrow\mathrm{ EXTZ((RA))
divisor }\leftarrow\mathrm{ EXTZ((RB))
remainder }\leftarrow\mathrm{ dividend % divisor
RT }\leftarrow\mathrm{ Chop(remainder, 64)

```

Let dividend be the unsigned integer doubleword in register RA.

Let divisor be the unsigned integer doubleword in register RB.

The remainder of dividend divided by divisor is placed into register RT. The quotient is not supplied as a result.

The remainder is the unique unsigned integer that satisfies
```

remainder = dividend . (quotient x divisor)

```
where \(0 \leq\) remainder < divisor.
If an attempt is made to perform any of the divisions
\[
\text { <anything> \% } 0
\]
then the contents of register RT are undefined.

\section*{Special Registers Altered: \\ None}

\subsection*{3.3.10 Fixed-Point Compare Instructions}

The fixed-point Compare instructions compare the contents of register RA with (1) the sign-extended value of the SI field, (2) the zero-extended value of the UI field, or (3) the contents of register RB. The comparison is signed for cmpi and cmp, and unsigned for cmpli and cmpl.

The \(L\) field controls whether the operands are treated as 64-bit or 32-bit quantities, as follows:

L Operand length
0 32-bit operands
1 64-bit operands
When the operands are treated as 32 -bit signed quantities, bit 32 of the register (RA or RB) is the sign bit.

The Compare instructions set one bit in the leftmost three bits of the designated CR field to 1 , and the other two to 0. XER \(_{\text {So }}\) is copied to bit 3 of the designated CR field.

The CR field is set as follows

Bit Name Description
\begin{tabular}{lll}
0 & LT & \begin{tabular}{l}
\((R A)<\) SI or (RB) (signed comparison) \\
\((R A)<\) U UI or (RB) (unsigned comparison)
\end{tabular} \\
1 & GT & \begin{tabular}{l} 
(RA) \(>\) SI or (RB) (signed comparison) \\
\((R A)>\mathrm{u}\) UI or (RB) (unsigned comparison)
\end{tabular} \\
2 & EQ & \begin{tabular}{l} 
(RA) \(=\) SI, UI, or (RB)
\end{tabular} \\
3 & SO & \begin{tabular}{l} 
Summary Overflow from the XER
\end{tabular}
\end{tabular}

\section*{Extended mnemonics for compares}

A set of extended mnemonics is provided so that compares can be coded with the operand length as part of the mnemonic rather than as a numeric operand. Some of these are shown as examples with the Compare instructions. See Appendix \(C\) for additional extended mnemonics.
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{Compare Immediate} & \multirow[t]{2}{*}{D-form} \\
\hline \multicolumn{6}{|l|}{cmpi BF,L,RA,SI} & \\
\hline \(0 \quad 11\) & \({ }_{6} \mathrm{BF}\) & \begin{tabular}{l|l} 
/ \\
9 & L \\
\hline
\end{tabular} & \begin{tabular}{l|l|l} 
L & RA \\
10 & 11
\end{tabular} & 16 & SI & 31 \\
\hline \multicolumn{7}{|l|}{\[
\text { if } \begin{aligned}
L= & 0 \text { then } a \leftarrow \operatorname{EXTS}\left((R A)_{32: 63)}\right) \\
& \text { else } a \leftarrow(R A)
\end{aligned}
\]} \\
\hline \multicolumn{7}{|l|}{if \(\quad \mathrm{a}<\operatorname{EXTS}(\mathrm{SI})\) then \(\mathrm{c} \leftarrow 0 \mathrm{~b} 100\)} \\
\hline \multicolumn{7}{|l|}{else if \(\mathrm{a}>\operatorname{EXTS}(\mathrm{SI})\) then \(\mathrm{c} \leftarrow 0 \mathrm{~b} 010\)} \\
\hline \multicolumn{7}{|l|}{else \(\quad c \leftarrow 0 \mathrm{~b} 001\)} \\
\hline \multicolumn{7}{|l|}{\(\mathrm{CR}_{4 \times \mathrm{BF}+32: 4 \times \mathrm{BF}+35} \leftarrow \mathrm{c} \| \mathrm{XER}_{\text {SO }}\)} \\
\hline
\end{tabular}

The contents of register RA ((RA) \({ }_{32: 63}\) sign-extended to 64 bits if \(\mathrm{L}=0\) ) are compared with the sign-extended value of the SI field, treating the operands as signed integers. The result of the comparison is placed into CR field BF.

\section*{Special Registers Altered:}
CR field BF

Extended Mnemonics:
Examples of extended mnemonics for Compare Immediate:
\begin{tabular}{ll} 
Extended: \\
cmpdi & \(R x\), value \\
cmpwi & \(c r 3, R x\), value
\end{tabular}

Equivalent to:
cmpi \(0,1, R x\),value
cmpi \(3,0, R x\),value


The contents of register RA \(\left((R A)_{32: 63}\right.\) if \(\left.L=0\right)\) are compared with the contents of register RB ((RB) 32:63 \(^{2}\) if \(\mathrm{L}=0\) ), treating the operands as signed integers. The result of the comparison is placed into CR field BF.
Special Registers Altered: CR field BF

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Compare:
\begin{tabular}{llll}
\multicolumn{2}{l}{ Extended: } & \multicolumn{2}{l}{ Equivalent to: } \\
cmpd & \(R x, R y\) & \(c m p\) & \(0,1, R x, R y\) \\
\(c m p w\) & \(c r 3, R x, R y\) & \(c m p\) & \(3,0, R x, R y\)
\end{tabular}

\section*{Version 3.0}

\section*{Compare Logical Immediate \\ D-form}
cmpli BF,L,RA,UI
\begin{tabular}{|l|l|l|l|l|l|lll|}
\hline 10 & & BF & & L & RA & & UI & \\
\hline 0 & & 6 & & 9 & 10 & 11 & 16 & \\
\hline
\end{tabular}


The contents of register RA ((RA) \({ }_{32: 63}\) zero-extended to 64 bits if \(L=0\) ) are compared with \({ }^{48} 0\) II UI, treating the operands as unsigned integers. The result of the comparison is placed into CR field BF.

\section*{Special Registers Altered:}

\section*{CR field BF}

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Compare Logical Immediate:

\section*{Extended:}
cmpldi Rx,value cmplwi cr3,Rx,value

\section*{Equivalent to:}
cmpli 0,1,Rx,value cmpli \(3,0, R x\), value

Compare Logical
X-form
cmpl BF,L,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline 31 & & BF & & L & RA & RB & & 32 \\
0 & & 6 & & 9 & 10 & 11 & 16 & 21 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { else } a \leftarrow \text { (RA) } \\
& b \leftarrow(\mathrm{RB}) \\
& \text { if } \quad a<{ }^{\mathrm{u}} \mathrm{~b} \text { then } \mathrm{c} \leftarrow 0 \mathrm{~b} 100 \\
& \text { else if } \mathrm{a}>{ }^{\mathrm{u}} \mathrm{~b} \text { then } \mathrm{c} \leftarrow 0 \mathrm{~b} 010 \\
& \text { else } \quad \mathrm{c} \leftarrow 0 \mathrm{~b} 001 \\
& \mathrm{CR}_{4 \times \mathrm{BF}+32: 4 \times \mathrm{BF}+35} \leftarrow \mathrm{C} \| \mathrm{XER}_{\mathrm{SO}}
\end{aligned}
\]

The contents of register RA ( \((\mathrm{RA})_{32: 63}\) if \(\mathrm{L}=0\) ) are compared with the contents of register RB ((RB) \()_{32: 63}\) if \(\mathrm{L}=0\) ), treating the operands as unsigned integers. The result of the comparison is placed into CR field BF.

\section*{Special Registers Altered:}

\section*{CR field BF}

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Compare Logical:
\begin{tabular}{llll} 
Extended: & \multicolumn{2}{l}{ Equivalent to: } \\
cmpld & \(R x, R y\) & \(c m p l\) & \(0,1, R x, R y\) \\
cmplw & \(c r 3, R x, R y\) & \(c m p l\) & \(3,0, R x, R y\)
\end{tabular}

\subsection*{3.3.10.1 Character-Type Compare Instructions}
```

Compare Ranged Byte X-form
cmprb BF,L,RA,RB

```

```

$\operatorname{src1} \leftarrow \operatorname{EXTZ}(\mid \text { RA })_{56: 63}$
stc21hi $\leftarrow$ ExTz( $(\text { RB })_{32: 39)}$
$\operatorname{sic} 2110 \leftarrow E X T Z(\mid \text { RB })_{40: 47)}$
$\operatorname{src22hi} \leftarrow$ EXTZ ( $(\text { RB })_{48: 55)}$
sic2210 $\leftarrow$ EXTZ $\left.(\mid R B)_{56: 63)}\right)$
if $\mathrm{L}=0$ then
in_range $\leftarrow(\operatorname{sic} 2210 \leq \operatorname{src} 1) \&(\operatorname{sic} 1 \leq \operatorname{src} 22 h i)$
else
in_range $\leftarrow((\operatorname{src} 2110 \leq \operatorname{src} 1) \&(\operatorname{sic} 1 \leq \operatorname{src} 21 h i)) \mid$
$\|(\operatorname{src} 22 \mid 0 \leq \operatorname{src}) \&(\operatorname{src} 1 \leq \operatorname{src} 22 h i) \mid$
$\mathrm{CR}_{4 \times 8 \mathrm{~B}+32} \leftarrow \mathrm{ObO}$
$\mathrm{CR}_{4 \times 8 \mathrm{~F}+33} \leftarrow$ in_range
$\mathrm{CR}_{4 \times 8 \mathrm{~B}+34} \leftarrow \mathrm{ObO}$
$\mathrm{CR}_{4 \times 8 \mathrm{~F}+35} \leftarrow \mathrm{ObO}$

```

Let srcl be the unsigned integer value in bits 56:63 of register RA.

Let src21hi be the unsigned integer value in bits 32:39 of register RB.

Let \(\operatorname{src} 2110\) be the unsigned integer value in bits 40:47 of register RB .

Let src22hi be the unsigned integer value in bits 48:55 of register RB.

Let src2210 be the unsigned integer value in bits 56:63 of register RB.

Let \(x\) be considered "in range" of \(y: z\) if the value \(x\) is greater than or equal to the value \(y\) and the value \(x\) is less than or equal to the value \(z\).

When \(L=0\), the value in_range is set to 1 if \(\operatorname{srcl}\) is in range of src2210:srci2hi. Otherwise, the value in_range is set to 0 .

When \(L=1\), the value in range is set to 1 if either srcl is in range of \(\operatorname{src} 21 \mid 0: 5 r c 21 \mathrm{hi}\), or srcl is in range of \(\operatorname{src} 2210: \operatorname{src} 22 h i\). Otherwise, the value in_range is set to 0 .
\(C R\) field \(B F\) is set to the value \(0 b 0\) concatenated with in_range concatenated with 0b00.

\section*{Special Registers Altered:}

CR field BF

\section*{- Programming Note}
cmprb is useful for implementing character typing functions such as isalpha(), isdigit(), isupper(), and islower() that are implemented using one or two range compares of the character.

A single-range compare can be implemented with an addi to load the upper and lower bounds in the range, such asisdigit().
\begin{tabular}{|c|c|c|}
\hline \(1 i\) & rRNG, \(0,0 \times 3930\) & loads ASCll values for 'g' and 'O' intor rRN \\
\hline cmpr b & cr TGT, O, rCHAR, rRNG & perform range compare sets CR field TGT to indicate in range \\
\hline
\end{tabular}

A combination of addi-addis can be used to set up 2 ranges, such as for isalpha().
```

Ii rRNG,O, Ox5A41 ; loads ASCII values for 'Z'
and 'A' into rRNG
lis rRNG,RNG, Ox7A61 ; appends ASCII values for 'z'
and 'a' into rRNG
cmprb crTGT,1,rCHAR,rRNG ; perform range compare on
character in rCHAR,
setting CR field TGT to
indicate in range

```

\section*{Version 3.0}

\section*{Compare Equal Byte X-form}
cmpeqb
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \begin{tabular}{l} 
BF
\end{tabular} & \begin{tabular}{l} 
I/ \\
9
\end{tabular} & RA \\
\hline
\end{tabular}
```

srcl \leftarrowGPR[RA],bit[56:63]
match}\leftarrow(srcl=(RB\mp@subsup{)}{00:07}{\prime}
(srcl = (RB) 08:15)
(srcl = (RB) 16:23)
(srCl =(RB) 24:31)
(srCl = (RB) 32:30)
(srcl = (RB) 40:47)
(srCl =(RB) 48:55)
(srcl = (RB) 56:63)
CR 4\times8F+32}<< ObO
CR4\times8F+33}\leftarrow~\mathrm{ match
CR4\timesBF+34}\leftarrow\textrm{ObO
CR4\times8F+35}\leftarrow~Ob

```
\(C R\) field \(B F\) is set to indicate if the contents of bits \(56: 63\) of register RA are equal to the contents of any of the 8 bytes in register RB.

Results are undefined in 32-bit mode.

\section*{Special Registers Altered:}
\(C R\) field \(B F\)

\section*{Programming Note}
cmpeqb is useful for implementing character typing functions such as isspace() that are implemented by comparing the character to 1 or more values.

A function such as isspace() can be implemented by loading the 6 byte codes corresponding to characters considered as whitespace (HT, LF, VT, FF, CR, and SP) and using the cmpeb to compare the subject character to those 6 values to determine if any match occurs.
```

Idx rSPC,WS_CHARS ; rSPC = OxOgOg_OgOA_OBOC_OD2O
load rSPC with all 6 ASCII
values corresponding to
white spaces
cmpeqb 2,crl,rCHAR,rSPC ; perform match compare on
character in rCHAR with
byte values in rSPC

```

In this case, the byte code for \(\mathrm{HT}(0 \times 09)\) was replicated to fill the all 8 bytes to avoid a potential miscompare.

\subsection*{3.3.11 Fixed-Point Trap Instructions}

The Trap instructions are provided to test for a specified set of conditions. If any of the conditions tested by a Trap instruction are met, the system trap handler is invoked. If none of the tested conditions are met, instruction execution continues normally.
The contents of register RA are compared with either the sign-extended value of the SI field or the contents of register RB, depending on the Trap instruction. For \(\boldsymbol{t d i}\) and \(\boldsymbol{t d}\), the entire contents of RA (and RB) participate in the comparison; for twi and \(t w\), only the contents of the low-order 32 bits of RA (and RB) participate in the comparison.
This comparison results in five conditions which are ANDed with TO. If the result is not 0 the system trap handler is invoked. These conditions are as follows.
\begin{tabular}{ll} 
TO Bit & ANDed with Condition \\
0 & Less Than, using signed comparison \\
1 & Greater Than, using signed comparison \\
2 & Equal \\
3 & Less Than, using unsigned comparison \\
4 & Greater Than, using unsigned comparison
\end{tabular}

\section*{Extended mnemonics for traps}

A set of extended mnemonics is provided so that traps can be coded with the condition as part of the mnemonic rather than as a numeric operand. Some of these are shown as examples with the Trap instructions. See Appendix C for additional extended mnemonics.

```

a}\leftarrow\operatorname{EXTS}((RA) 32:63
if (a < EXTS(SI)) \& TO
if (a > EXTS(SI)) \& TO
if (a = EXTS(SI)) \& TO2 then TRAP
if (a<u}\operatorname{EXTS}(SI))\& TO then TRAP
if (a>政 EXTS(SI)) \& TO

```

The contents of \(R A_{32: 63}\) are compared with the sign-extended value of the SI field. If any bit in the TO field is set to 1 and its corresponding condition is met by the result of the comparison, the system trap handler is invoked.

If the trap conditions are met, this instruction is context synchronizing (see Book III).
Special Registers Altered: None

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Trap Word Immediate:
\begin{tabular}{lllr}
\multicolumn{2}{l}{ Extended: } & \multicolumn{2}{l}{ Equivalent to: } \\
twgti & \(R x\), value & twi & \(8, R x\),value \\
twllei & \(R x\), value & twi & \(6, R x\), value
\end{tabular}

```

a \leftarrow EXTS((RA) 32:63)
b}\leftarrow\operatorname{EXTS}((\textrm{RB}\mp@subsup{)}{32:63)}{
if (a<b) \& TO0 then TRAP
if (a>b) \& TO
if (a=b) \& TO2 then TRAP
if (a<u b) \& TOO
if (a > }\mp@subsup{}{}{u}\mathrm{ b) \& TO4 then TRAP

```

The contents of \(\mathrm{RA}_{32: 63}\) are compared with the contents of \(\mathrm{RB}_{32: 63 \text {. If any bit in the TO field is set to } 1 \text { and }}\) its corresponding condition is met by the result of the comparison, the system trap handler is invoked.

If the trap conditions are met, this instruction is context synchronizing (see Book III).

\section*{Special Registers Altered:}

None

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Trap Word:
\begin{tabular}{llll}
\multicolumn{2}{c}{ Extended: } & \multicolumn{2}{c}{ Equivalent to: } \\
tweq & \(R x, R y\) & tw & \(4, R x, R y\) \\
twlge & \(R x, R y\) & tw & \(5, R x, R y\) \\
trap & & tw & \(31,0,0\)
\end{tabular}

\subsection*{3.3.11.1 64-bit Fixed-Point Trap Instructions}

\section*{Trap Doubleword Immediate D-form}
tdi TO,RA,S
\begin{tabular}{|l|l|l|lll|}
\hline 2 & TO & RA & & SI & 31 \\
\hline 0 & 6 & 11 & 16 & & 3 \\
\hline
\end{tabular}
```

a}\leftarrow(RA
b}\leftarrow\operatorname{EXTS}(SI
if (a<b) \& TOO then TRAP
if (a>b) \& TO
if (a = b) \& TO2 then TRAP
if (a<u}b)\& TO_ then TRAP
if (a > }\mp@subsup{}{}{\textrm{u}}\mathrm{ b) \& TO

```

The contents of register RA are compared with the sign-extended value of the SI field. If any bit in the TO field is set to 1 and its corresponding condition is met by the result of the comparison, the system trap handler is invoked.

If the trap conditions are met, this instruction is context synchronizing (see Book III).

\section*{Special Registers Altered:}

None
Extended Mnemonics:
Examples of extended mnemonics for Trap Doubleword Immediate:
\begin{tabular}{lllr}
\multicolumn{2}{l}{ Extended: } & \multicolumn{2}{l}{ Equivalent to: } \\
tdlti & \(R x\),value & tdi & \(16, R x\), value \\
tdnei & \(R x\), value & tdi & \(24, R x\), value
\end{tabular}

\subsection*{3.3.12 Fixed-Point Select}

\section*{Integer Select}

A-form
isel \(\quad R T, R A, R B, B C\)
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & BC & 15 & 1 \\
0 & & 6 & 11 & 16 & 21 & 26 \\
31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { if } \mathrm{RA}=0 \text { then } a \leftarrow 0 \text { else } a \leftarrow \text { (RA) } \\
& \text { if } \mathrm{CR}_{\mathrm{BC}+32}=1 \text { then } \mathrm{RT} \leftarrow a \\
& \text { else } \mathrm{RT} \leftarrow(\mathrm{RB})
\end{aligned}
\]

If the contents of bit \(\mathrm{BC}+32\) of the Condition Register are equal to 1 , then the contents of register RA (or 0 ) are placed into register RT. Otherwise, the contents of register RB are placed into register RT.

\section*{Special Registers Altered:}

None

Trap Doubleword
X-form

```

a\leftarrow(RA)
b}\leftarrow(\textrm{RB}
if (a<b) \& TOO then TRAP
if (a>b)\&TO
if (a = b)\& TO2 then TRAP
if (a<u}b)\&TO\mp@code{Then TRAP
if (a>u

```

The contents of register RA are compared with the contents of register RB. If any bit in the TO field is set to 1 and its corresponding condition is met by the result of the comparison, the system trap handler is invoked.
If the trap conditions are met, this instruction is context synchronizing (see Book III).
Special Registers Altered:
None

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Trap Doubleword:
\begin{tabular}{ll} 
Extended: & Equivalent to: \\
tdge \(R x, R y\) & td \(\quad 12, R x, R y\)
\end{tabular}

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Integer Select.
\begin{tabular}{ll}
\multicolumn{2}{c}{ Extended: } \\
isellt & \(R x, R y, R z\) \\
iselgt & \(R x, R y, R z\) \\
iseleq & \(R x, R y, R z\)
\end{tabular}

Equivalent to:
isel \(R x, R y, R z, 0\)
isel \(R x, R y, R z, 1\)
isel \(R x, R y, R z, 1\)

\subsection*{3.3.13 Fixed-Point Logical Instructions}

The Logical instructions perform bit-parallel operations on 64-bit operands.

The X-form Logical instructions with \(\mathrm{Rc}=1\), and the D-form Logical instructions andi. and andis., set the first three bits of CR Field 0 as described in Section 3.3.8, "Other Fixed-Point Instructions" on page 67. The Logical instructions do not change the I SO, OV, OV32, CA, and CA32 bits in the XER.

\section*{Extended mnemonics for logical operations}

Extended mnemonics are provided that generate two different types of "no-ops" (instructions that do nothing). The first type is the preferred form, which is optimized to minimize its use of the processor's execution resources. This form is based on the OR Immediate instruction. The second type is the executed form, which is intended to consume the same amount of the processor's execution resources as if it were not a
no-op. This form is based on the XOR Immediate instruction. (There are also no-ops that have other uses, such as affecting program priority, for which extended mnemonics have not been defined.)

Extended mnemonics are provided that use the \(O R\) and \(N O R\) instructions to copy the contents of one register to another, with and without complementing. These are shown as examples with the two instructions.
See Appendix C, "Assembler Extended Mnemonics" on page 795 for additional extended mnemonics.

\section*{Programming Note}

Warning: Some forms of no-op may have side effects such as affecting program priority. Programmers should use the preferred no-op unless the side effects of some other form of no-op are intended.

AND Immediate
D-form
andi. RA,RS,UI
\begin{tabular}{|c|c|c|ccc|}
\hline 28 & RS & RA & \multicolumn{2}{|c|}{} & UI \\
\hline 0 & & 6 & 11 & 16 & \\
\hline
\end{tabular}
\(R A \leftarrow(R S) \&\left({ }^{48} 0 \| U I\right)\)
The contents of register RS are ANDed with \({ }^{48} 0 \mathrm{II} \mathrm{UI}\) and the result is placed into register RA.
Special Registers Altered:
CRO
AND Immediate Shifted
D-form
andis. RA,RS,UI
\begin{tabular}{|c|c|c|ccc|}
\hline 29 & RS & RA & \multicolumn{2}{|c|}{ Ul } & \\
\hline 0 & & 6 & 11 & 16 & \\
\hline
\end{tabular}
\[
\text { RA } \leftarrow(\text { RS }) \&\left({ }^{32} 0 \| \text { UI } \|{ }^{16} 0\right)
\]

The contents of register RS are ANDed with \({ }^{32} 0\) II UI II \({ }^{16} 0\) and the result is placed into register RA.

\section*{Special Registers Altered:}

CRO

OR Immediate
D-form
ori RA,RS,UI
\begin{tabular}{|l|l|l|lll|}
\hline 24 & RS & RA & \multicolumn{2}{|c|}{} & UI \\
\hline 0 & & 6 & 11 & 16 & \\
\hline
\end{tabular}
\(\mathrm{RA} \leftarrow(\mathrm{RS}) \mid\left({ }^{48} 0| | \mathrm{UI}\right)\)
The contents of register RS are ORed with \({ }^{48} \mathrm{OIIUI}\) and the result is placed into register RA.

The preferred "no-op" (an instruction that does nothing) is:
\[
\text { ori } \quad 0,0,0
\]

Special Registers Altered: None

\section*{Extended Mnemonics:}

Example of extended mnemonics for OR Immediate:
\begin{tabular}{ll} 
Extended: & \multicolumn{2}{l}{ Equivalent to: } \\
no-op & ori \(0,0,0\)
\end{tabular}

\section*{Version 3.0}

\section*{OR Immediate Shifted}

D-form
oris RA,RS,UI
\begin{tabular}{|l|l|l|lll|}
\hline 25 & \multicolumn{2}{|c|}{ RS } & RA & \multicolumn{2}{|c|}{ Ul } \\
\hline 0 & 6 & 11 & 16 & & 31 \\
\hline
\end{tabular}
\(\mathrm{RA} \leftarrow(\mathrm{RS}) \mid\left({ }^{32} 0| | \mathrm{UI}| |^{16} 0\right)\)
The contents of register RS are ORed with \({ }^{32} 0\) II UI II \({ }^{16} 0\) and the result is placed into register RA.

\section*{Special Registers Altered:}

None

\section*{XOR Immediate}

\section*{D-form}
xori RA,RS,UI
\begin{tabular}{|l|l|l|lll|}
\hline 26 & RS & RA & & UI & 31 \\
\hline 0 & & 6 & 11 & 16 & \\
\hline
\end{tabular}
\[
\mathrm{RA} \leftarrow(\mathrm{RS}) \text { XOR }\left({ }^{48} 0 \| \mathrm{UI}\right)
\]

The contents of register RS are XORed with \({ }^{48} 0\) II UI and the result is placed into register RA.

The executed form of a "no-op" (an instruction that does nothing, but consumes execution resources nevertheless) is:
\[
\text { xori } 0,0,0
\]

\section*{Special Registers Altered:}

None

\section*{Extended Mnemonics:}

Example of extended mnemonics for XOR Immediate:
\begin{tabular}{ll} 
Extended: & \multicolumn{1}{l}{ Equivalent to: } \\
xnop & xori \(0,0,0\)
\end{tabular}

\section*{Programming Note}

The executed form of no-op should be used only when the intent is to alter the timing of a program.

\(R A \leftarrow(R S) \&(R B)\)
The contents of register RS are ANDed with the contents of register RB and the result is placed into register RA.

Some forms of and Rx, Rx, Rx provide special functions; see Section 9.3 of Book III.

\section*{Special Registers Altered: \\ CRO}
(if \(R c=1\) )
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{5}{|l|}{XOR} & \multicolumn{2}{|r|}{X-form} \\
\hline xor & \multicolumn{4}{|c|}{RA,RS,RB} & \multicolumn{2}{|r|}{(Rc=0)} \\
\hline xor. & \multicolumn{4}{|c|}{RA,RS,RB} & \multicolumn{2}{|r|}{(Rc=1)} \\
\hline 31 & RS & RA & RB & & 316 & Rc \\
\hline 0 & 6 & 11 & 16 & 21 & & 31 \\
\hline
\end{tabular}
\(\mathrm{RA} \leftarrow(\mathrm{RS}) \oplus(\mathrm{RB})\)
The contents of register RS are XORed with the contents of register RB and the result is placed into register RA.

\section*{Special Registers Altered:}

CRO
(if \(\mathrm{Rc}=1\) )

\section*{NAND}
\begin{tabular}{lll} 
nand & \(R A, R S, R B\) & \((R c=0)\) \\
nand. & \(R A, R S, R B\) & \((R c=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 31 & RS & & RA & RB & & 476 & Rc \\
0 & & 6 & 11 & 16 & 21 & & 31 \\
\hline
\end{tabular}
```

RA}\leftarrow\neg((RS)\& (RB)

```

The contents of register RS are ANDed with the contents of register RB and the complemented result is placed into register RA.

\section*{Special Registers Altered:}

\section*{CR0}
(if \(\mathrm{Rc}=1\) )

\section*{Programming Note}
nand or nor with RS=RB can be used to obtain the one's complement.

```

RA}\leftarrow(RS
(RB)

```

The contents of register RS are ORed with the contents of register RB and the result is placed into register RA.

Some forms of or \(R x, R x, R x\) provide special functions; see Section 3.2 and Section 4.3.3, both in Book II.

Special Registers Altered:
CRO
(if \(\mathrm{Rc}=1\) )
Extended Mnemonics:
Example of extended mnemonics for OR:
\begin{tabular}{ll} 
Extended: & Equivalent to: \\
\(m r \quad R x, R y\) & or \(\quad R x, R y, R y\)
\end{tabular}


The contents of register RS are ORed with the contents of register RB and the complemented result is placed into register RA.

\section*{Special Registers Altered: CRO \\ (if \(\mathrm{Rc}=1\) )}

\section*{Extended Mnemonics:}

Example of extended mnemonics for NOR:
\begin{tabular}{ll} 
Extended: & Equivalent to: \\
not \(\quad R x, R y\) & nor \(\quad R x, R y, R y\)
\end{tabular}

\section*{AND with Complement}

X-form
\begin{tabular}{lll} 
andc & \(R A, R S, R B\) & \((R c=0)\) \\
andc. & \(R A, R S, R B\) & \((R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|cc|c|}
\hline 31 & RS & RA & RB & & 60 & Rc \\
0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}

\section*{\(R A \leftarrow(R S) \& \neg(R B)\)}

The contents of register RS are ANDed with the complement of the contents of register RB and the result is placed into register RA.

\section*{Special Registers Altered:}
(if \(\mathrm{Rc}=1\) )

Equivalent
\(X\)-form

```

RA}\leftarrow(\textrm{RS})\equiv(\textrm{RB}

```

The contents of register RS are XORed with the contents of register RB and the complemented result is placed into register RA.

\section*{Special Registers Altered:}

CRO
(if \(\mathrm{Rc}=1\) )

OR with Complement
X-form
\begin{tabular}{lll} 
orc & \(R A, R S, R B\) & \((R c=0)\) \\
orc. & \(R A, R S, R B\) & \((R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 412 \\
\hline 0 & & 11 & 16 & 21 & \\
\hline 6 \\
\hline
\end{tabular}
\(R A \leftarrow(R S) \mid \neg(R B)\)
The contents of register RS are ORed with the complement of the contents of register RB and the result is placed into register RA.
Special Registers Altered:
CRO
(if \(\mathrm{Rc}=1\) )
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{5}{|l|}{Extend Sign Byte} & \multicolumn{2}{|r|}{X-form} \\
\hline extsb & \multicolumn{4}{|l|}{RA,RS} & \multicolumn{2}{|r|}{(Rc=0)} \\
\hline extsb. & RA, & & & & & =1) \\
\hline 31 & RS & RA & I/I & & 954 & Rc \\
\hline 0 & 6 & 11 & 16 & 21 & & 31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& s \leftarrow(R S)_{56} \\
& R A_{56: 63} \leftarrow(\mathrm{RS})_{56: 63} \\
& R A_{0: 55} \leftarrow 56_{\mathrm{S}}
\end{aligned}
\]
\((\mathrm{RS})_{56: 63}\) are placed into \(\mathrm{RA}_{56: 63} \cdot \mathrm{RA}_{0: 55}\) are filled with a copy of (RS) 56 .
Special Registers Altered:
CRO
(if \(\mathrm{Rc}=1\) )
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{5}{|l|}{Extend Sign Halfword} & \multicolumn{2}{|r|}{\(X\)-form} \\
\hline extsh & \multicolumn{4}{|l|}{RA,RS} & \multicolumn{2}{|r|}{(Rc=0)} \\
\hline extsh. & RA, & & & & & =1) \\
\hline 31 & RS & RA & //] & & 922 & Rc \\
\hline 0 & 6 & 11 & 16 & 21 & & 31 \\
\hline
\end{tabular}
\(S \leftarrow(R S) 48\)
\(R A_{48: 63} \leftarrow(R S) 48: 63\)
\(R A_{0: 47} \leftarrow 48{ }_{S}\)
\((\mathrm{RS})_{48: 63}\) are placed into \(\mathrm{RA}_{48: 63} . \mathrm{RA}_{0: 47}\) are filled with a copy of (RS) 48 .
Special Registers Altered:
CRO
(if \(\mathrm{Rc}=1\) )
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{5}{|l|}{Count Leading Zeros Word} & \multicolumn{2}{|r|}{\multirow[t]{2}{*}{X-form
(Rc=0)}} \\
\hline cntlizw
cntlzw. & RA, & \[
\begin{aligned}
& \mathrm{S} \\
& (\mathrm{Rc}=1)
\end{aligned}
\] & & & & \\
\hline 31 & RS & 11 RA & 16 & 21 & 26 & Rc \\
\hline
\end{tabular}
```

n}\leftarrow3
do while n < 64
if (RS)}\mp@subsup{n}{n}{}=1\mathrm{ then leave
n}\leftarrow\textrm{n}+
RA}\leftarrow\textrm{n}-3

```

A count of the number of consecutive zero bits starting at bit 32 of register RS is placed into register RA. This number ranges from 0 to 32, inclusive.

If \(R c\) is equal to \(1, C R\) field 0 is set to reflect the result.

\section*{Special Registers Altered:} CRO
(if \(\mathrm{Rc}=1\) )
Programming Note
For both Count Leading Zeros instructions, if Rc=1 then LT is set to 0 in CR Field 0.

\section*{Count Trailing Zeros Word}
\begin{tabular}{lll} 
cnttzw & RA, RS & \((R C=0)\) \\
cnttzw. & \(R A, R S\) & \((R C=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 31 & & RS & & RA & & III & \\
\hline 0 & 538 & & \\
\hline 31
\end{tabular}
\(\mathrm{n} \leftarrow 0\)
do while n < 32
if (RS) \(63-\mathrm{n}=0 \mathrm{~b} 1\) then leave
\(\mathrm{n} \leftarrow \mathrm{n}+1\)
\(R A \leftarrow \operatorname{EXTZ64(n)}\)
A count of the number of consecutive zero bits starting at bit 63 of the rightmost word of register RS is placed into register RA. This number ranges from 0 to 32, inclusive.

If \(R C\) is equal to \(1, C R\) field 0 is set to reflect the result.
Special Registers Altered:
CRO
(if \(R C=1\) )

\section*{Version 3.0}

\section*{Compare Bytes}

\section*{X-form}
cmpb
RA,RS,RB

do \(\mathrm{n}=0\) to 7
\[
\begin{aligned}
& \text { if } \mathrm{RS}_{8 \times \mathrm{n}: 8 \times \mathrm{n}+7}=(\mathrm{RB})_{8 \times n: 8 \times \mathrm{n}+7} \text { then } \\
& \mathrm{RA}_{8 \times \mathrm{n}: 8 \times \mathrm{n}+7} \leftarrow{ }^{8} 1 \\
& \text { else } \\
& \quad \mathrm{RA}_{8 \times \mathrm{n}: 8 \times \mathrm{n}+7} \leftarrow{ }^{8} 0
\end{aligned}
\]

Each byte of the contents of register RS is compared to each corresponding byte of the contents in register RB. If they are equal, the corresponding byte in RA is set to 0xFF. Otherwise the corresponding byte in RA is set to \(0 \times 00\).

\section*{Special Registers Altered:}

None

Population Count Bytes
X-form
```

popcntb RA, RS

```
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & I/I & & 122 & \(/\) \\
0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

do $i=0$ to 7
$n \leftarrow 0$
do $j=0$ to 7
if $(\mathrm{RS})_{(i \times 8)+j}=1$ then
$\mathrm{n} \leftarrow \mathrm{n}+1$
$\mathrm{RA}_{(\mathrm{i} \times 8):(\mathrm{i} \times 8)+7} \leftarrow \mathrm{n}$

```

A count of the number of one bits in each byte of register RS is placed into the corresponding byte of register RA. This number ranges from 0 to 8 , inclusive.

\section*{Special Registers Altered:}

None

\section*{Population Count Words \\ X-form}
```

popcntw RA, RS

```
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & \({ }^{\prime} / / I\) & & 378 & \\
0 & & 6 & & 11 & 16 & 21 \\
31 \\
\hline
\end{tabular}
```

do i = 0 to 1
n}\leftarrow
do j = 0 to 31
if (RS)}\mp@subsup{}{(i\times32)+j}{}=1\mathrm{ then
n}\leftarrow\textrm{n}+
RA

```

A count of the number of one bits in each word of register RS is placed into the corresponding word of register RA. This number ranges from 0 to 32, inclusive.

\section*{Special Registers Altered:}

None
Parity Doubleword
prtyd RA,RS
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & RA & & I/I & \\
\hline 0 & & 6 & 186 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

S}\leftarrow
do i = 0 to 7
S}\leftarrow\textrm{S}\oplus(\textrm{RS}\mp@subsup{)}{i\times8+7}{
RA}\leftarrow\mp@subsup{}{}{63}0||

```

The least significant bit in each byte of the contents of register RS is examined. If there is an odd number of one bits the value 1 is placed into register RA; otherwise the value 0 is placed into register RA.

Special Registers Altered:
None

Parity Word X-form
prtyw RA,RS
\begin{tabular}{|l|l|l|l|ll|c|}
\hline 31 & \multicolumn{1}{|c|}{RS} & \multicolumn{2}{c|}{ RA } & \multicolumn{1}{c|}{\(/ / /\)} & & 154 \\
\hline 0 & & 6 & & 11 & 16 & 21 \\
\hline
\end{tabular}
```

s}\leftarrow
t\leftarrow0
do i = 0 to 3
s}\leftarrow\textrm{S}\oplus(\textrm{RS}\mp@subsup{)}{i\times8+7}{
do i = 4 to 7
t}\leftarrowt\oplus(RS\mp@subsup{)}{i\times8+7}{
RA 0:31}\leftarrow\mp@subsup{}{}{31}0||\textrm{S
RA 32:63}\leftarrow\mp@subsup{}{}{31}0||

```

The least significant bit in each byte of \((R S)_{0: 31}\) is examined. If there is an odd number of one bits the value 1 is placed into \(R A_{0: 31}\); otherwise the value 0 is placed into \(\mathrm{RA}_{0: 31}\). The least significant bit in each byte of \((\mathrm{RS})_{32: 63}\) is examined. If there is an odd number of one bits the value 1 is placed into \(\mathrm{RA}_{32: 63}\); otherwise the value 0 is placed into \(R A_{32: 63}\).
Special Registers Altered:
None

\section*{Programming Note}

The Parity instructions are designed to be used in conjunction with the Population Count instruction to compute the parity of words or a doubleword. The parity of the upper and lower words in (RS) can be computed as follows.
popentb RA, RS
prtyw RA, RA
The parity of (RS) can be computed as follows.
```

popentb RA, RS
prtyd RA, RA

```

\subsection*{3.3.13.1 64-bit Fixed-Point Logical Instructions}


\section*{Count Leading Zeros Doubleword X-form}
\begin{tabular}{lll} 
cntIzd & RA,RS & \((\mathrm{Rc}=0)\) \\
cntlzd. & \(\mathrm{RA}, \mathrm{RS}\) & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|}
\hline & 31 & & RS & & RA & & III & & 58 & Rc \\
\hline & & 6 & & 11 & & 16 & & 21 & & 31 \\
\hline
\end{tabular}
```

n}\leftarrow
do while n < 64
if (RS)}n=1\mathrm{ then leave
n}\leftarrow\textrm{n}+
RA}\leftarrow\textrm{n

```

A count of the number of consecutive zero bits starting at bit 0 of register RS is placed into register RA. This number ranges from 0 to 64, inclusive.

If \(\mathrm{Rc}=1, \mathrm{CR}\) Field 0 is set to reflect the result.
Special Registers Altered:

\section*{CRO}
(if \(\mathrm{Rc}=1\) )

\section*{Count Trailing Zeros Doubleword}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline cnttzd & \multicolumn{4}{|c|}{RA,RS} & \multicolumn{2}{|r|}{( \(R C=0\) )} \\
\hline cnttzd. & & A,RS & & & & \\
\hline \[
\begin{array}{ll} 
& 31 \\
0
\end{array}
\] & \[
\sigma_{6} \mathrm{RS}
\] & \[
{ }_{11} R A
\] & \[
{ }_{16} \quad \text { III }
\] & 21 & 570 & \(|\)\begin{tabular}{|r|}
\(R C\) \\
31
\end{tabular} \\
\hline
\end{tabular}
```

n}\leftarrow
do while n < 64
if (RS) 63-n = Ob1 then leave
n}\leftarrow\textrm{n}+
RA \leftarrow EXTZ64(n)

```

A count of the number of consecutive zero bits starting at bit 63 of register RS is placed into register RA. This number ranges from 0 to 64, inclusive.

If \(R C\) is equal to \(1, C R\) field 0 is set to reflect the result.
Special Registers Altered:
CRO

\section*{Bit Permute Doubleword X-form}
bpermd RA,RS,RB]
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 252 & 1 \\
\hline 0 & & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

For i = 0 to 7
index }\leftarrow(\textrm{RS}\mp@subsup{)}{8*i:8*i+7}{*
If index < 64

```

```

        else permi
    RA}\leftarrow\mp@subsup{}{}{56}0|| perm0:7

```

Eight permuted bits are produced. For each permuted bit i where i ranges from 0 to 7 and for each byte \(i\) of RS, do the following.

If byte \(i\) of \(R S\) is less than 64, permuted bit \(i\) is set to the bit of RB specified by byte i of RS; otherwise permuted bit i is set to 0 .

The permuted bits are placed in the least-significant byte of RA, and the remaining bits are filled with Os.

\section*{Special Registers Altered:}

None

\section*{Programming Note}

The fact that the permuted bit is 0 if the corresponding index value exceeds 63 permits the permuted bits to be selected from a 128-bit quantity, using a single index register. For example, assume that the 128 -bit quantity \(Q\), from which the permuted bits are to be selected, is in registers r2 (high-order 64 bits of Q) and r3 (low-order 64 bits of \(Q\) ), that the index values are in register r1, with each byte of \(r 1\) containing a value in the range \(0: 127\), and that each byte of register r4 contains the value 64. The following code sequence selects eight permuted bits from \(Q\) and places them into the low-order byte of r 6 .
bpermd r6,r1,r2
\# select from highorder half of Q
xor \(r 0, r 1, r 4\) \# adjust index values
bpermd r5,r0,r3 \# select from low-
order half of \(Q\)
or r6,r6,r5
\# merge the two selections

\subsection*{3.3.14 Fixed-Point Rotate and Shift Instructions}

The Fixed-Point Facility performs rotation operations on data from a GPR and returns the result, or a portion of the result, to a GPR.

The rotation operations rotate a 64-bit quantity left by a specified number of bit positions. Bits that exit from position 0 enter at position 63.

Two types of rotation operation are supported.
For the first type, denoted rotate \(_{64}\) or ROTL \(_{64}\), the value rotated is the given 64 -bit value. The rotate 64 operation is used to rotate a given 64-bit quantity.

For the second type, denoted rotate \({ }_{32}\) or \(\mathrm{ROTL}_{32}\), the value rotated consists of two copies of bits 32:63 of the given 64-bit value, one copy in bits 0:31 and the other in bits \(32: 63\). The rotate 32 operation is used to rotate a given 32-bit quantity.
The Rotate and Shift instructions employ a mask generator. The mask is 64 bits long, and consists of 1 -bits from a start bit, mstart, through and including a stop bit, mstop, and 0-bits elsewhere. The values of mstart and mstop range from 0 to 63 . If mstart > mstop, the 1 -bits wrap around from position 63 to position 0 . Thus the mask is formed as follows:
\[
\begin{aligned}
& \text { if mstart } \leq \text { mstop then } \\
& \text { mask }_{\text {mstart }: \text { mstop }}=\text { ones } \\
& \text { mask }_{\text {all other bits }}=\text { zeros } \\
& \text { else } \\
& \text { mask }_{\text {mstart: } 63}=\text { ones } \\
& \text { mask }_{0: \text { mstop }}=\text { ones } \\
& \text { mask }_{\text {all other bits }}=\text { zeros }
\end{aligned}
\]

There is no way to specify an all-zero mask.
For instructions that use the rotate 32 operation, the mask start and stop positions are always in the low-order 32 bits of the mask.
The use of the mask is described in following sections.
The Rotate and Shift instructions with Rc=1 set the first three bits of CR field 0 as described in Section 3.3.8, "Other Fixed-Point Instructions" on page 67. Rotate and I Shift instructions do not change the OV, OV32, and SO bits. Rotate and Shift instructions, except algebraic right shifts, do not change the CA and CA32 bits.

\section*{Extended mnemonics for rotates and shifts}

The Rotate and Shift instructions, while powerful, can be complicated to code (they have up to five operands). A set of extended mnemonics is provided that allow simpler coding of often-used functions such as clearing the leftmost or rightmost bits of a register, left justifying or right justifying an arbitrary field, and performing simple rotates and shifts. Some of these are shown as examples with the Rotate instructions. See Appendix C, "Assembler Extended Mnemonics" on page 795 for additional extended mnemonics.

\subsection*{3.3.14.1 Fixed-Point Rotate Instructions}

These instructions rotate the contents of a register. The result of the rotation is

■ inserted into the target register under control of a mask (if a mask bit is 1 the associated bit of the rotated data is placed into the target register, and if the mask bit is 0 the associated bit in the target register remains unchanged); or
- ANDed with a mask before being placed into the target register.
The Rotate Left instructions allow right-rotation of the contents of a register to be performed (in concept) by a left-rotation of \(64-\mathrm{n}\), where n is the number of bits by which to rotate right. They allow right-rotation of the contents of the low-order 32 bits of a register to be performed (in concept) by a left-rotation of 32-n, where \(n\) is the number of bits by which to rotate right.

\section*{Rotate Left Word Immediate then AND with Mask M-form}
\begin{tabular}{lll} 
rlwinm & RA, RS, SH,MB,ME & \((\mathrm{Rc}=0)\) \\
rlwinm. & \(\mathrm{RA}, \mathrm{RS}, \mathrm{SH}, \mathrm{MB}, \mathrm{ME}\) & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 21 & RS & \multicolumn{1}{|c|}{ RA } & SH & MB & ME & Rc \\
0 & & 6 & 11 & 16 & 21 & 26 \\
\hline
\end{tabular}
```

n}\leftarrowS
r}\leftarrow\mp@subsup{\textrm{ROTL}}{32}{}((\textrm{RS})32:63,\textrm{n}
m}\leftarrow\operatorname{MASK}(MB+32, ME+32
RA}\leftarrowr\&

```

The contents of register RS are rotated 32 left SH bits. A mask is generated having 1-bits from bit MB+32 through bit \(\mathrm{ME}+32\) and 0 -bits elsewhere. The rotated data are ANDed with the generated mask and the result is placed into register RA.
Special Registers Altered: CRO (if Rc=1)

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Rotate Left Word Immediate then AND with Mask:
\begin{tabular}{ll}
\multicolumn{2}{l}{ Extended: } \\
extlwi & \(R x, R y, n, b\) \\
srwi & \(R x, R y, n\) \\
clrrwi & \(R x, R y, n\)
\end{tabular}

\section*{Equivalent to:}
rlwinm Rx,Ry,b,0,n-1
rlwinm Rx,Ry,32-n,n,31
rlwinm Rx,Ry,0,0,31-n

\section*{Programming Note}

Let RSL represent the low-order 32 bits of register RS, with the bits numbered from 0 through 31 .
rlwinm can be used to extract an n-bit field that starts at bit position b in RSL, right-justified into the low-order 32 bits of register RA (clearing the remaining \(32-n\) bits of the low-order 32 bits of RA), by setting \(S H=b+n, M B=32-n\), and \(M E=31\). It can be used to extract an n-bit field that starts at bit position b in RSL, left-justified into the low-order 32 bits of register RA (clearing the remaining 32-n bits of the low-order 32 bits of RA), by setting \(\mathrm{SH}=\mathrm{b}\), \(M B=0\), and \(M E=n-1\). It can be used to rotate the contents of the low-order 32 bits of a register left (right) by \(n\) bits, by setting \(\mathrm{SH}=\mathrm{n}(32-\mathrm{n}), \mathrm{MB}=0\), and \(\mathrm{ME}=31\). It can be used to shift the contents of the low-order 32 bits of a register right by \(n\) bits, by setting \(\mathrm{SH}=32-\mathrm{n}, \mathrm{MB}=\mathrm{n}\), and \(\mathrm{ME}=31\). It can be used to clear the high-order \(b\) bits of the low-order 32 bits of the contents of a register and then shift the result left by \(n\) bits, by setting \(S H=n, M B=b-n\), and \(M E=31-n\). It can be used to clear the low-order \(n\) bits of the low-order 32 bits of a register, by setting \(\mathrm{SH}=0, \mathrm{MB}=0\), and \(\mathrm{ME}=31-\mathrm{n}\).
For all the uses given above, the high-order 32 bits of register RA are cleared.

Extended mnemonics are provided for all of these uses; see Appendix C, "Assembler Extended Mnemonics" on page 795.

\title{
Rotate Left Word then AND with Mask M-form
}
\begin{tabular}{lll} 
rlwnm & RA, RS, RB, MB,ME & \((R c=0)\) \\
rlwnm. & \(R A, R S, R B, M B, M E\) & \((R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 23 & RS & RA & RB & MB & & Rc \\
\hline \(\bigcirc\) & & & 6 & & & \\
\hline
\end{tabular}
```

n}\leftarrow(\textrm{RB}\mp@subsup{)}{59:63}{
r}\leftarrow\mp@subsup{\operatorname{ROTL}}{32}{((RS)
m}\leftarrowMASK (MB+32, ME+32
RA}\leftarrowr\&

```

The contents of register RS are rotated \({ }_{32}\) left the number of bits specified by \((\mathrm{RB})_{59: 63}\). A mask is generated having 1-bits from bit MB+32 through bit ME+32 and 0 -bits elsewhere. The rotated data are ANDed with the generated mask and the result is placed into register RA.

\section*{Special Registers Altered:}
CR0
(if \(\mathrm{Rc}=1\) )

\section*{Extended Mnemonics:}

Example of extended mnemonics for Rotate Left Word then AND with Mask:
\begin{tabular}{ll} 
Extended: & Equivalent to: \\
rotlw \(\quad R x, R y, R z\) & rlwnm \(\quad R x, R y, R z, 0,31\)
\end{tabular}

Extended:
rotlw Rx,Ry,Rz
rlwnm Rx,Ry,Rz,0,31

\section*{Programming Note}

Let RSL represent the low-order 32 bits of register RS, with the bits numbered from 0 through 31.
rlwnm can be used to extract an n-bit field that starts at variable bit position b in RSL, right-justified into the low-order 32 bits of register RA (clearing the remaining 32-n bits of the low-order 32 bits of \(R A\) ), by setting \(R_{59: 63}=b+n, \quad M B=32-n\), and \(\mathrm{ME}=31\). It can be used to extract an \(n\)-bit field that starts at variable bit position b in RSL, left-justified into the low-order 32 bits of register RA (clearing the remaining \(32-n\) bits of the low-order 32 bits of \(R A\) ), by setting \(R_{59: 63}=b, M B=0\), and \(M E=n-1\). It can be used to rotate the contents of the low-order 32 bits of a register left (right) by variable \(n\) bits, by setting \(\mathrm{RB}_{59: 63}=\mathrm{n}(32-\mathrm{n}), \mathrm{MB}=0\), and \(\mathrm{ME}=31\).
For all the uses given above, the high-order 32 bits of register RA are cleared.

Extended mnemonics are provided for some of these uses; see Appendix C, "Assembler Extended Mnemonics" on page 795.

\section*{Rotate Left Word Immediate then Mask Insert M-form}
\begin{tabular}{lll} 
rlwimi & \(R A, R S, S H, M B, M E\) & \((R c=0)\) \\
rlwimi. & \(R A, R S, S H, M B, M E\) & \((R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 20 & RS & \multicolumn{1}{|c|}{ RA } & SH & MB & ME & Rc \\
0 & & 6 & 11 & 16 & 21 & 26 \\
\hline
\end{tabular}
```

n}\leftarrow\textrm{SH
r}\leftarrow\mp@subsup{\operatorname{ROTL}}{32}{((RS) 32:63,n)
m}\leftarrowMASK(MB+32,ME+32
RA}\leftarrowr\&m|(RA)\&\urcorner

```

The contents of register RS are rotated \({ }_{32}\) left SH bits. A mask is generated having 1-bits from bit MB+32 through bit \(\mathrm{ME}+32\) and 0 -bits elsewhere. The rotated data are inserted into register RA under control of the generated mask.

\section*{Special Registers Altered:}
CRO
(if \(\mathrm{Rc}=1\) )

\section*{Extended Mnemonics:}

Example of extended mnemonics for Rotate Left Word Immediate then Mask Insert.

\section*{Extended:}

\section*{Equivalent to:}
inslwi Rx,Ry,n,b
rlwimi Rx,Ry,32-b,b,b+n-1

\section*{Programming Note}

Let RAL represent the low-order 32 bits of register RA, with the bits numbered from 0 through 31.
rlwimi can be used to insert an n-bit field that is left-justified in the low-order 32 bits of register RS, into RAL starting at bit position b, by setting \(S H=32-b, M B=b\), and \(M E=(b+n)-1\). It can be used to insert an n-bit field that is right-justified in the low-order 32 bits of register RS, into RAL starting at bit position \(b\), by setting \(S H=32-(b+n), M B=b\), and \(M E=(b+n)-1\).

Extended mnemonics are provided for both of these uses; see Appendix C, "Assembler Extended Mnemonics" on page 795.

\subsection*{3.3.14.1.1 64-bit Fixed-Point Rotate Instructions}

\section*{Rotate Left Doubleword Immediate then Clear Left MD-form}
\begin{tabular}{lll} 
rldicl & \(\mathrm{RA}, \mathrm{RS}, \mathrm{SH}, \mathrm{MB}\) & \((\mathrm{Rc}=0)\) \\
rldicl. & \(\mathrm{RA}, \mathrm{RS}, \mathrm{SH}, \mathrm{MB}\) & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 30 & RS & \multicolumn{1}{|c|}{ RA } & \multicolumn{1}{|c|}{ sh } & mb & 0 & sh & Rc \\
0 & & 6 & 11 & 16 & 21 & 27 & 30 \\
\hline
\end{tabular}
```

n}\leftarrow\mp@subsup{\textrm{sh}}{5}{}||\mp@subsup{\textrm{sh}}{0:4}{
r}\leftarrow\mp@subsup{\textrm{ROTL}}{64}{((RS), n)
b}\leftarrow\mp@subsup{\textrm{mb}}{5}{|| m\mp@subsup{b}{0:4}{}
m}\leftarrowMASK (b, 63
RA}\leftarrowr\&

```

The contents of register RS are rotated \({ }_{64}\) left SH bits. A mask is generated having 1-bits from bit MB through bit 63 and 0 -bits elsewhere. The rotated data are ANDed with the generated mask and the result is placed into register RA.

\section*{Special Registers Altered:}
CRO
(if \(\mathrm{Rc}=1\) )

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Rotate Left Doubleword Immediate then Clear Left.
\begin{tabular}{ll}
\multicolumn{2}{l}{ Extended: } \\
extrdi & Rx,Ry,n,b \\
srdi & Rx,Ry,n \\
clrldi & Rx,Ry,n
\end{tabular}

\section*{Equivalent to:}
\begin{tabular}{ll} 
rldicl & \(R x, R y, b+n, 64-n\) \\
rldicl & \(R x, R y, 64-n, n\) \\
rldicl & \(R x, R y, 0, n\)
\end{tabular}

\section*{Programming Note}
rldicl can be used to extract an n-bit field that starts at bit position \(b\) in register RS, right-justified into register RA (clearing the remaining 64-n bits of \(R A\) ), by setting \(S H=b+n\) and \(M B=64-n\). It can be used to rotate the contents of a register left (right) by \(n\) bits, by setting \(S H=n(64-n)\) and \(M B=0\). It can be used to shift the contents of a register right by n bits, by setting \(\mathrm{SH}=64-\mathrm{n}\) and \(\mathrm{MB}=\mathrm{n}\). It can be used to clear the high-order \(n\) bits of a register, by setting \(\mathrm{SH}=0\) and \(\mathrm{MB}=\mathrm{n}\).

Extended mnemonics are provided for all of these uses; see Appendix C, "Assembler Extended Mnemonics" on page 795.

\title{
Rotate Left Doubleword Immediate then Clear Right MD-form
}
\begin{tabular}{lll} 
rldicr & RA,RS,SH,ME & \((\mathrm{Rc}=0)\) \\
rldicr. & \(\mathrm{RA}, \mathrm{RS}, \mathrm{SH}, \mathrm{ME}\) & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 30 & RS & RA & \multicolumn{1}{|c|}{ sh } & me & 1 & sh & Rc \\
\hline 0 & & 6 & 11 & 16 & 21 & 27 & 30 \\
\hline
\end{tabular}
```

n}\leftarrow\mp@subsup{\textrm{sh}}{5}{}||s\mp@subsup{h}{0:4}{
r}\leftarrow\mp@subsup{\textrm{ROTL}}{64}{((RS), n)
e}\leftarrowm\mp@subsup{e}{5}{}||m\mp@subsup{e}{0:4}{
m}\leftarrow\operatorname{MASK}(0,e
RA}\leftarrowr\&

```

The contents of register RS are rotated \({ }_{64}\) left SH bits. A mask is generated having 1 -bits from bit 0 through bit ME and 0-bits elsewhere. The rotated data are ANDed with the generated mask and the result is placed into register RA.

\section*{Special Registers Altered:}
CRO
(if \(\mathrm{Rc}=1\) )

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Rotate Left Doubleword Immediate then Clear Right.
\begin{tabular}{ll}
\multicolumn{2}{l}{ Extended: } \\
extldi & \(R x, R y, n, b\) \\
sldi & \(R x, R y, n\) \\
clrrdi & \(R x, R y, n\)
\end{tabular}

Equivalent to:
\begin{tabular}{ll} 
rldicr & \(R x, R y, b, n-1\) \\
rldicr & \(R x, R y, n, 63-n\) \\
rldicr & \(R x, R y, 0,63-n\)
\end{tabular}

\section*{Programming Note}
rldicr can be used to extract an \(n\)-bit field that starts at bit position \(b\) in register RS, left-justified into register RA (clearing the remaining 64-n bits of RA), by setting \(S H=b\) and \(M E=n-1\). It can be used to rotate the contents of a register left (right) by n bits, by setting \(\mathrm{SH}=\mathrm{n}\) ( \(64-\mathrm{n}\) ) and \(\mathrm{ME}=63\). It can be used to shift the contents of a register left by \(n\) bits, by setting \(S H=n\) and \(M E=63-n\). It can be used to clear the low-order \(n\) bits of a register, by setting \(\mathrm{SH}=0\) and \(\mathrm{ME}=63-\mathrm{n}\).
Extended mnemonics are provided for all of these uses (some devolve to rldicl); see Appendix C, "Assembler Extended Mnemonics" on page 795.

\section*{Rotate Left Doubleword Immediate then Clear MD-form}
\begin{tabular}{lll} 
rldic & \(R A, R S, S H, M B\) & \((R c=0)\) \\
rldic. & \(R A, R S, S H, M B\) & \((R c=1)\)
\end{tabular}

\[
\begin{aligned}
& \mathrm{n} \leftarrow \mathrm{sh}_{5} \| \mathrm{sh}_{0: 4} \\
& \mathrm{r} \leftarrow \mathrm{ROTL}_{64}((\mathrm{RS}), \mathrm{n}) \\
& \mathrm{b} \leftarrow \mathrm{mb}_{5} \| \mid \mathrm{mb}_{0}: 4 \\
& \mathrm{~m} \leftarrow \operatorname{MASK}(\mathrm{~b}, \mathrm{n}) \\
& \mathrm{RA} \leftarrow \mathrm{r} \& \mathrm{~m}
\end{aligned}
\]

The contents of register RS are rotated \({ }_{64}\) left SH bits. A mask is generated having 1 -bits from bit MB through bit 63-SH and 0-bits elsewhere. The rotated data are ANDed with the generated mask and the result is placed into register RA.

\section*{Special Registers Altered:}

\section*{CR0}
(if \(\mathrm{Rc}=1\) )

\section*{Extended Mnemonics:}

Example of extended mnemonics for Rotate Left Doubleword Immediate then Clear:

\section*{Extended:}
clrisldi Rx,Ry,b,n

\section*{Equivalent to:}
\[
0
\]

Programming Note
rldic can be used to clear the high-order \(b\) bits of the contents of a register and then shift the result left by \(n\) bits, by setting \(\mathrm{SH}=\mathrm{n}\) and \(\mathrm{MB}=\mathrm{b}-\mathrm{n}\). It can be used to clear the high-order \(n\) bits of a register, by setting \(\mathrm{SH}=0\) and \(\mathrm{MB}=\mathrm{n}\).
Extended mnemonics are provided for both of these uses (the second devolves to rldicl); see Appendix C, "Assembler Extended Mnemonics" on page 795.

\section*{Rotate Left Doubleword then Clear Left MDS-form}
\begin{tabular}{lll} 
rldcl & RA, RS, RB, MB & (Rc=0) \\
rldcl. & \(R A, R S, R B, M B\) & \((R c=1)\)
\end{tabular}

\[
\begin{aligned}
& n \leftarrow(\mathrm{RB})_{58: 63} \\
& r \leftarrow \mathrm{ROTL}_{64}((\mathrm{RS}), \mathrm{n}) \\
& \mathrm{b} \leftarrow \mathrm{mb}_{5}| | \mathrm{mb}_{0}: 4 \\
& m \leftarrow M A S K(\mathrm{~b}, 63) \\
& \mathrm{RA} \leftarrow \mathrm{r} \& \mathrm{~m}
\end{aligned}
\]

The contents of register RS are rotated \({ }_{64}\) left the number of bits specified by (RB) \(58: 63\). A mask is generated having 1-bits from bit MB through bit 63 and 0 -bits elsewhere. The rotated data are ANDed with the generated mask and the result is placed into register RA.

\section*{Special Registers Altered:}

\section*{CRO}
(if \(R c=1\) )

\section*{Extended Mnemonics:}

Example of extended mnemonics for Rotate Left Doubleword then Clear Left.
\begin{tabular}{ll} 
Extended: & Equivalent to: \\
rotld \(\quad R x, R y, R z\) & rldcl \(\quad R x, R y, R z, 0\)
\end{tabular}

\section*{Programming Note}
rldcl can be used to extract an n-bit field that starts at variable bit position \(b\) in register RS, right-justified into register RA (clearing the remaining 64-n bits of RA), by setting \(\mathrm{RB}_{58: 63}=b+n\) and \(M B=64-n\). It can be used to rotate the contents of a register left (right) by variable n bits, by setting \(\mathrm{RB}_{58: 63}=\mathrm{n}\) ( \(64-n\) ) and \(M B=0\).
Extended mnemonics are provided for some of these uses; see Appendix C, "Assembler Extended Mnemonics" on page 795.

\section*{Rotate Left Doubleword then Clear Right MDS-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline rldcr & \multicolumn{4}{|l|}{\multirow[t]{2}{*}{RA,RS,RB,ME RA,RS,RB,ME}} & \multicolumn{2}{|l|}{\multirow[t]{2}{*}{\[
\begin{aligned}
& (\mathrm{Rc}=0) \\
& (\mathrm{Rc}=1)
\end{aligned}
\]}} \\
\hline rldcr. & & & & & & \\
\hline 30 & RS & RA & RB & me & 9 & Rc \\
\hline 0 & & & & & & 31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \mathrm{n} \leftarrow(\mathrm{RB})_{58: 63} \\
& \mathrm{r} \leftarrow \mathrm{ROTL}_{64}((\mathrm{RS}), \mathrm{n}) \\
& \mathrm{e} \leftarrow \mathrm{me}_{5}| | \mathrm{me}_{0: 4} \\
& \mathrm{~m} \leftarrow \operatorname{MASK}(0, \mathrm{e}) \\
& \mathrm{RA} \leftarrow \mathrm{r} \& \mathrm{~m}
\end{aligned}
\]

The contents of register RS are rotated \({ }_{64}\) left the number of bits specified by (RB) \({ }_{58: 63}\). A mask is generated having 1-bits from bit 0 through bit ME and 0-bits elsewhere. The rotated data are ANDed with the generated mask and the result is placed into register RA.

\section*{Special Registers Altered:}

\section*{CR0}
(if \(\mathrm{Rc}=1\) )

\section*{Programming Note}
rldcr can be used to extract an n -bit field that starts at variable bit position \(b\) in register RS, left-justified into register RA (clearing the remaining 64-n bits of \(R A\) ), by setting \(R B_{58: 63}=b\) and \(M E=n-1\). It can be used to rotate the contents of a register left (right) by variable n bits, by setting \(\mathrm{RB}_{58: 63}=\mathrm{n}\) ( \(64-n\) ) and ME=63.

Extended mnemonics are provided for some of these uses (some devolve to rldcl); see Appendix C, "Assembler Extended Mnemonics" on page 795.

\section*{Rotate Left Doubleword Immediate then Mask Insert MD-form}
\begin{tabular}{lll} 
rldimi & \(R A, R S, S H, M B\) & \((R c=0)\) \\
rldimi. & \(R A, R S, S H, M B\) & \((R c=1)\)
\end{tabular}

```

n}\leftarrow\mp@subsup{\textrm{sh}}{5}{}||\mp@subsup{\textrm{sh}}{0:4}{4
r}\leftarrow\mp@subsup{\textrm{ROTL}}{64}{((RS), n)
b}\leftarrow\mp@subsup{\textrm{mb}}{5}{|| mb 0:4
m}\leftarrow\operatorname{MASK}(\textrm{b}, ᄀn
RA \leftarrowr\&m | (RA)\& \m

```

The contents of register RS are rotated \({ }_{64}\) left SH bits. A mask is generated having 1-bits from bit MB through bit \(63-\mathrm{SH}\) and 0 -bits elsewhere. The rotated data are inserted into register RA under control of the generated mask.
Special Registers Altered:
CRO
(if \(\mathrm{Rc}=1\) )

\section*{Extended Mnemonics:}

Example of extended mnemonics for Rotate Left Doubleword Immediate then Mask Insert.

\section*{Extended:}
insrdi Rx,Ry,n,b

Equivalent to:
rldimi \(R x, R y, 64-(b+n), b\)

\section*{Programming Note}
rldimi can be used to insert an n-bit field that is right-justified in register RS, into register RA starting at bit position \(b\), by setting \(\mathrm{SH}=64-(\mathrm{b}+\mathrm{n})\) and \(M B=b\).

An extended mnemonic is provided for this use; see Appendix C, "Assembler Extended Mnemonics" on page 795.

\subsection*{3.3.14.2 Fixed-Point Shift Instructions}

The instructions in this section perform left and right shifts.

\section*{Extended mnemonics for shifts}

Immediate-form logical (unsigned) shift operations are obtained by specifying appropriate masks and shift values for certain Rotate instructions. A set of extended mnemonics is provided to make coding of such shifts simpler and easier to understand. Some of these are shown as examples with the Rotate instructions. See Appendix C, "Assembler Extended Mnemonics" on page 795 for additional extended mnemonics.

\section*{Programming Note}

Any Shift Right Algebraic instruction, followed by addze, can be used to divide quickly by \(2^{n}\). The setting of the CA and CA32 bits by the Shift Right Algebraic instructions is independent of mode.

\section*{Programming Note}

Multiple-precision shifts can be programmed as shown in Section E.1, "Multiple-Precision Shifts" on page 639.
```

n}\leftarrow(RB)59:6
r}\leftarrow\mp@subsup{\textrm{ROTL}}{32}{}((\textrm{RS}\mp@subsup{)}{32:63,n)}{n
if (RB) 58 = 0 then
m}\leftarrow\operatorname{MASK}(32,63-n
else m}\leftarrow\mp@subsup{}{}{64}
RA}\leftarrowr\&

```

The contents of the low-order 32 bits of register RS are shifted left the number of bits specified by (RB) \(58: 63\). Bits shifted out of position 32 are lost. Zeros are supplied to the vacated positions on the right. The 32-bit result is placed into \(\mathrm{RA}_{32: 63} . \mathrm{RA}_{0: 31}\) are set to zero. Shift amounts from 32 to 63 give a zero result.
Special Registers Altered:
(if \(R c=1\) )
Shift Right Word X-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline srw & \multicolumn{2}{|l|}{\multirow[t]{2}{*}{\[
\begin{aligned}
& \text { RA,RS,RB } \\
& \text { RA,RS,RB }
\end{aligned}
\]}} & & & \multicolumn{2}{|r|}{\multirow[t]{2}{*}{\((\mathrm{Rc}=0\)
\((\mathrm{Rc}=1)\)}} \\
\hline srw. & & & & & & \\
\hline 31 & RS & RA & RB & & 536 & Rc \\
\hline 0 & & 11 & 16 & 21 & & 31 \\
\hline
\end{tabular}
```

n}\leftarrow(\textrm{RB}\mp@subsup{)}{59:63}{

```

```

if (RB)58 = 0 then
m}\leftarrowMASK(n+32, 63
else m}\leftarrow\mp@subsup{}{}{64}
RA}\leftarrowr\&

```

The contents of the low-order 32 bits of register RS are shifted right the number of bits specified by \((\mathrm{RB})_{58: 63}\). Bits shifted out of position 63 are lost. Zeros are supplied to the vacated positions on the left. The 32-bit result is placed into \(\mathrm{RA}_{32: 63} . \mathrm{RA}_{0: 31}\) are set to zero. Shift amounts from 32 to 63 give a zero result.
Special Registers Altered:
CRO
(if \(\mathrm{Rc}=1\) )

\section*{Shift Right Algebraic Word Immediate \(X\)-form}
\begin{tabular}{lll} 
srawi & \(R A, R S, S H\) & \((R c=0)\) \\
srawi. & \(R A, R S, S H\) & \((R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & \({ }_{6} \mathrm{RS}\) & \({ }_{11} \mathrm{RA}\) & \({ }_{16} \mathrm{SH}\) & & 824 \\
\hline 0 & & & & & \\
31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \mathrm{n} \leftarrow \mathrm{SH} \\
& r \leftarrow \operatorname{ROTL}_{32}\left((\mathrm{RS})_{32: 63}, 64-n\right) \\
& m \leftarrow \operatorname{MASK}(\mathrm{n}+32,63) \\
& s \leftarrow(\mathrm{RS})_{32} \\
& \left.\mathrm{RA} \leftarrow \mathrm{r} \& \mathrm{~m} \mid\left({ }^{64} \mathrm{~S}\right) \&\right\urcorner \mathrm{m} \\
& \text { carry } \leftarrow s \&\left((r \& \neg m)_{32: 63} \neq 0\right) \\
& \text { CA } \leftarrow \text { carry } \\
& \text { CA32 } \leftarrow \text { carry }
\end{aligned}
\]

The contents of the low-order 32 bits of register RS are shifted right SH bits. Bits shifted out of position 63 are lost. Bit 32 of RS is replicated to fill the vacated positions on the left. The 32-bit result is placed into \(\mathrm{RA}_{32: 63}\). Bit 32 of RS is replicated to fill \(\mathrm{RA}_{0: 31}\). CA and CA32 are set to 1 if the low-order 32 bits of (RS) contain a negative number and any 1 -bits are shifted out of position 63; otherwise CA and CA32 are set to 0 . A shift amount of zero causes RA to receive EXTS((RS) 32:63 \(^{2}\) ), I and CA and CA32 to be set to 0

\section*{Special Registers Altered:}

I CA CA32
CRO
(if \(\mathrm{Rc}=1\) )
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{5}{|l|}{Shift Right Algebraic Word} & \multicolumn{2}{|r|}{X-form} \\
\hline sraw & \multicolumn{4}{|l|}{RA,RS,RB} & \multicolumn{2}{|r|}{(Rc=0)} \\
\hline sraw. & & RS,RB & & & & c=1) \\
\hline 31 & RS & RA & RB & & 792 & Rc \\
\hline 0 & 6 & 11 & 16 & 21 & & 31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \mathrm{n} \leftarrow(\mathrm{RB})_{59: 63} \\
& r \leftarrow \text { ROTL }_{32}\left((\mathrm{RS})_{32: 63}, 64-\mathrm{n}\right) \\
& \text { if }(\mathrm{RB})_{58}=0 \text { then } \\
& m \leftarrow \operatorname{MASK}(n+32,63) \\
& \text { else } m \leftarrow{ }^{64} 0 \\
& \mathrm{~s} \leftarrow(\mathrm{RS})_{32} \\
& \left.R A \leftarrow r \& m \mid\left({ }^{64} s\right) \&\right\urcorner m \\
& \text { carry } \leftarrow \mathrm{s} \&\left((\mathrm{r} \& \neg \mathrm{~m})_{32: 63} \neq 0\right) \\
& \mathrm{CA} \leftarrow \text { carry } \\
& \text { CA32 } \leftarrow \text { carry }
\end{aligned}
\]

The contents of the low-order 32 bits of register RS are shifted right the number of bits specified by \((\mathrm{RB})_{58: 63}\). Bits shifted out of position 63 are lost. Bit 32 of RS is replicated to fill the vacated positions on the left. The 32-bit result is placed into \(\mathrm{RA}_{32: 63}\). Bit 32 of RS is repli-
I cated to fill \(\mathrm{RA}_{0: 31}\). CA and CA32 are set to 1 if the low-order 32 bits of (RS) contain a negative number and any 1 -bits are shifted out of position 63; otherwise CA and CA32 are set to 0 . A shift amount of zero causes RA to receive \(\operatorname{EXTS}\left((\mathrm{RS})_{32: 63}\right)\), and CA and CA32 to be set to 0 . Shift amounts from 32 to 63 give a result of 64 sign bits, and cause CA and CA32 to receive the sign bit of \((\mathrm{RS})_{32: 63}\)

\section*{Special Registers Altered:}
\| CA CA32
CRO
(if Rc=1)

\subsection*{3.3.14.2.1 64-bit Fixed-Point Shift Instructions}

\section*{Shift Left Doubleword}
\begin{tabular}{ll} 
sld & \(R A, R S, R B\) \\
sld. & \(R A, R S, R B\)
\end{tabular}

X-form
( \(\mathrm{Rc}=0\) )
( \(\mathrm{Rc}=1\) )
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 27 & Rc \\
0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

n}\leftarrow(\textrm{RB}\mp@subsup{)}{58:63}{
r}\leftarrow\mp@subsup{ROTL}{64}{((RS), n)
if (RB) 57 = 0 then
m}\leftarrow\operatorname{MASK}(0,63-n
else m}\leftarrow\mp@subsup{}{}{64}
RA}\leftarrowr\&

```

The contents of register RS are shifted left the number of bits specified by (RB) \()_{57: 63}\). Bits shifted out of position 0 are lost. Zeros are supplied to the vacated positions on the right. The result is placed into register RA. Shift amounts from 64 to 127 give a zero result.

\section*{Special Registers Altered:}

CRO
(if \(\mathrm{Rc}=1\) )

Shift Right Doubleword
X-form
srd RA,RS,RB (Rc=0)
srd. RA,RS,RB (Rc=1)
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 539 & Rc \\
\hline 0 & 6 & 11 & 16 & 21 & & 31 \\
\hline
\end{tabular}
```

n}\leftarrow(\textrm{RB}\mp@subsup{)}{58:63}{
r}\leftarrow\mp@subsup{\textrm{ROTL}}{64}{(}(\textrm{RS}), 64-n
if (RB)57 = 0 then
m}\leftarrow\operatorname{MASK(n, 63)
else m}\leftarrow\mp@subsup{}{}{64}
RA}\leftarrowr\&

```

The contents of register RS are shifted right the number of bits specified by \((\mathrm{RB})_{57: 63}\). Bits shifted out of position 63 are lost. Zeros are supplied to the vacated positions on the left. The result is placed into register RA. Shift amounts from 64 to 127 give a zero result.

\section*{Special Registers Altered:}

CRO
(if \(\mathrm{Rc}=1\) )

\section*{Shift Right Algebraic Doubleword}

```

$\mathrm{n} \leftarrow \mathrm{sh}_{5}| | \mathrm{sh}_{0: 4}$
$r \leftarrow$ ROTL $_{64}((R S), 64-n)$
$\mathrm{m} \leftarrow \operatorname{MASK}(\mathrm{n}, 63)$
$s \leftarrow(\mathrm{RS})_{0}$
$\left.R A \leftarrow r \& m \mid\left({ }^{64} s\right) \&\right\urcorner m$
carry $\leftarrow \mathrm{s} \&((\mathrm{r} \& \neg \mathrm{~m}) \neq 0)$
$\mathrm{CA} \leftarrow$ carry
CA32 $\leftarrow$ carry

```

The contents of register RS are shifted right SH bits. Bits shifted out of position 63 are lost. Bit 0 of RS is replicated to fill the vacated positions on the left. The result
| is placed into register RA. CA and CA32 are set to 1 if (RS) is negative and any 1-bits are shifted out of posi-
I tion 63; otherwise CA and CA32 are set to 0 . A shift amount of zero causes RA to be set equal to (RS), and
I CA and CA32 to be set to 0 .

\section*{Special Registers Altered:}
l CA CA32
CRO
(if \(\mathrm{Rc}=1\) )

\section*{Shift Right Algebraic Doubleword X-form}

\[
\begin{aligned}
& n \leftarrow(R B)_{58: 63} \\
& r \leftarrow R O T L_{64}((R S), 64-n) \\
& \text { if }(R B)_{57}=0 \text { then } \\
& \quad m \leftarrow M A S K(n, 63) \\
& \text { else } \mathrm{m} \leftarrow 640 \\
& S \leftarrow(R S)_{0} \\
& R A \leftarrow r \& m \mid\left({ }^{64} S\right) \& \neg m \\
& \text { Carry } \leftarrow \mathrm{S} \&((r \& \neg m) \neq 0) \\
& C A \leftarrow \operatorname{carry} \\
& \text { CA32 } \leftarrow \text { carry }
\end{aligned}
\]

The contents of register RS are shifted right the number of bits specified by \((\mathrm{RB})_{57: 63}\). Bits shifted out of position 63 are lost. Bit 0 of RS is replicated to fill the vacated positions on the left. The result is placed into register RA. CA and CA32 are set to 1 if (RS) is negative and any 1-bits are shifted out of position 63; otherwise CA and CA32 are set to 0 . A shift amount of zero causes RA to be set equal to (RS), and CA and CA32 to be set to 0 . Shift amounts from 64 to 127 give a - result of 64 sign bits in RA, and cause CA and CA32 to receive the sign bit of (RS).
Special Registers Altered:
I
CA CA32
CRO
(if \(\mathrm{Rc}=1\) )

Extend-Sign Word and Shift Left Immediate XS-form
\begin{tabular}{lll} 
extswsli & \(R A, R S, S H\) & \((R C=0)\) \\
extswsli. & \(R A, R S, S H\) & \((R C=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{2}{|c|}{ RS } & \multicolumn{2}{c|}{ RA } & sh & \\
\hline 0 & & 445 & shRC \\
3031 \\
\hline
\end{tabular}
```

n < Sh | | sho:4
r \&ROTL64(EXTS64(RS 32:63), n)
m}\leftarrowMASK(0, 63\cdotn
RA\leftarrowr\&m

```

The contents of the low order 32 bits of RS are sign-extended to 64 bits and then shifted left SH bits. Bits shifted out of bit 0 are lost. Zeros are supplied to vacated bits on the right. The result is placed in register RA.

\section*{Special Registers Altered:}

CRO
(if \(R c=1\) )

\subsection*{3.3.15 Binary Coded Decimal (BCD) Assist Instructions}

The Binary Coded Decimal Assist instructions operate on Binary Coded Decimal operands (cbcdtd and
addg6s) and Decimal Floating-Point operands (cdtbcd) See Chapter 5. for additional information.

\section*{Convert Declets To Binary Coded Decimal X-form}
cdtbcd RA, RS
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & RS & RA & \(/ / /\) & 282 & \begin{tabular}{l}
1 \\
31 \\
\hline 0
\end{tabular}
\end{tabular}
```

do i = 0 to 1
n}\leftarrow\textrm{i}\times3
RA m+0:n+7
RA n+8:n+19}\leftarrow~DPD_TO_BCD( (RS) n+12:n+21) (
RA A+20:n+31}\leftarrow DPD_TO_BCD((RS) (R+22:n+31),

```

The low-order 20 bits of each word of register RS contain two declets which are converted to six, 4-bit BCD fields; each set of six, 4-bit BCD fields is placed into the low-order 24 bits of the corresponding word in RA. The high-order 8 bits in each word of RA are set to 0 .

\section*{Special Registers Altered:} None

\section*{Convert Binary Coded Decimal To Declets X-form}
cbcdtd
RA, RS
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & RS & RA & \(1 / /\) & 314 & 1 \\
\hline 0 & & 6 & & 11 \\
\hline
\end{tabular}
```

do i = 0 to 1
n}\leftarrow\mathrm{ i x }3
RA
RA
RA A+22:n+31

```

The low-order 24 bits of each word of register RS contain six, 4-bit BCD fields which are converted to two declets; each set of two declets is placed into the low-order 20 bits of the corresponding word in RA. The high-order 12 bits in each word of RA are set to 0 .

If a 4-bit BCD field has a value greater than 9 the results are undefined.

\section*{Special Registers Altered:}

None

Add and Generate Sixes
XO-form addg6s RT,RA,RB
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & RT & RA & RB & 1 & 74 & 7 \\
21
\end{tabular}
```

do i = 0 to 15
dci
c}\leftarrow\mp@subsup{}{}{4}(\mp@subsup{\textrm{dc}}{0}{})|\mp@subsup{|}{}{4}(\mp@subsup{\textrm{dc}}{1}{})||···|\mp@subsup{|}{}{4}(\mp@subsup{\textrm{dc}}{15}{}
RT \leftarrow (\negC) \& 0x6666_6666_6666_6666

```

The contents of register RA are added to the contents of register RB. Sixteen carry bits are produced, one for each carry out of decimal position \(n\) (bit position \(4 x n\) ).
A doubleword is composed from the 16 carry bits, and placed into RT. The doubleword consists of a decimal six (0b0110) in every decimal digit position for which the corresponding carry bit is 0 , and a zero (0b0000) in every position for which the corresponding carry bit is 1.

\section*{Special Registers Altered:}

None

\section*{Programming Note}
addg6s can be used to add or subtract two BCD operands. In these examples it is assumed that r0 contains 0x666...666. (BCD data formats are described in Section 5.3.)

Addition of the unsigned BCD operand in register RA to the unsigned BCD operand in register RB can be accomplished as follows.
\begin{tabular}{ll} 
add & \(r 1, R A, r 0\) \\
add & \(r 2, r 1, R B\) \\
addg6s & \(R T, r 1, R B\) \\
subf & \(R T, R T, r 2 \# R T=R A+{ }_{B C D} R B\)
\end{tabular}

Subtraction of the unsigned BCD operand in register RA from the unsigned BCD operand in register RB can be accomplished as follows. (In this example it is assumed that RB is not register 0 .)
```

addi r1,RB,1
nor r2,RA,RA\# one's complement of RA
add r3,r1,r2
addg6s RT,r1,r2
subf RT,RT,r3\# RT = RB -

```

Additional instructions are needed to handle signed BCD operands, and BCD operands that occupy more than one register (e.g., unsigned BCD operands that have more than 16 decimal digits).

\subsection*{3.3.16 Move To/From Vector-Scalar Register Instructions}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline mfvsrd & & RA,XS & & & & \\
\hline \[
31
\] & \[
{ }_{6} \mathrm{~S}
\] & \[
{ }_{11} \mathrm{RA}
\] & \[
{ }_{16} \quad I I I
\] & 21 & 51 & Sx \\
\hline
\end{tabular}
if \(S X=0\) \& MSR. \(F P=0\) then \(F P\) _Unavailable()
if \(S X=1\) \& MSR.VEC=O then Véctor_Unavailablell
GPR[RA] \& VSR[ \(32 \times 5 X+S]\). dwor d[ 0\(]\)
Let \(X S\) be the value \(32 \times 5 X+S\).
The contents of doubleword element 0 of VSR[ XS] are placed into GPR[ RA].

For \(S X=0\), mfvsrd is treated as a Floating-Point instruction in terms of resource availability.

For \(S X=1\), \(\boldsymbol{m f v s r} \boldsymbol{d}\) is treated as a Vector instruction in terms of resource availability.
\begin{tabular}{llll}
\multicolumn{2}{l}{ Extended Mnemonics } & \multicolumn{2}{l}{ Equivalent To } \\
\(m f\) fprd & RA, FRS & mf vsrd & RA, FRS \\
\(m f v r d\) & RA, VRS & mfvsrd & RA, VRS +32
\end{tabular}

Special Registers Altered
None
Data Layout for mfvsrd

unused

\section*{Move From VSR Lower Doubleword XX1-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|l|}{mfvsrld RA,XS} \\
\hline \[
31
\] & \({ }_{6} \mathrm{~S}\) & \[
{ }_{11} \mathrm{RA}
\] & \[
{ }_{16} \quad I I \prime
\] & 21 & 307 & \(\left\lvert\, \begin{aligned} & 5 x \\ & 31\end{aligned}\right.\) \\
\hline
\end{tabular}
if \(S X=0\) \& MSR. VSX=0 then VSX_Unavailablell
if \(S X=1\) \& MSR.VEC=O then Vector_Unavailable el)
GPR[ RA] \(\leftarrow\) VSR[ \(32 \times 5 X+5]\). dwor d [1]
Let \(X S\) be the value \(32 \times 5 x+S\).
The contents of doubleword 1 of VSR[XS] are placed into GPR[RA] .

For \(S X=0\), mfvsrld is treated as a Floating-Point instruction in terms of resource availability.

For \(S X=1\), \(\boldsymbol{m} f\) vsrld is treated as a Vector instruction in terms of resource availability.

Special Registers Altered: None

Data Layout for mfvsrid
src = VSR[ XS]
\begin{tabular}{|l|l|}
\hline unused &. dword \([1]\) \\
\hline
\end{tabular}
tgt = GPR[RA]


\section*{Version 3.0}

\section*{Move From VSR Word and Zero XX1-form}

\section*{mfvsrwz RA,XS}
\begin{tabular}{|l|l|l|l|l|l|}
\hline 31 & & S & & RA & \multicolumn{1}{|c|}{ III } \\
0 & & 6 & & 115 & SX \\
\hline 11 & & 16 & & & \\
\hline
\end{tabular}
if \(S X=0\) \& MSR. \(F P=0\) then \(F P\) _Unavailablell
if \(S X=1\) \& MSR.VEC=O then Véctor_Unavailablel)
GPR[ RA] \(\leftarrow\) EXTZ64 [ (VSR[ \(32 \times 5 X+5]\). wor d [ 1\(]\) )
Let \(X S\) be the value \(32 \times 5 X+S\).
The contents of word element 1 of VSR[XS] are placed into bits 32:63 of GPR[RA]. The contents of bits 0:31 of GPR[RA] are set to 0 .

For \(S X=0\), mfvsrwz is treated as a Floating-Point instruction in terms of resource availability.

For \(S X=1, \boldsymbol{m} f \mathbf{v s r w z}\) is treated as a Vector instruction in terms of resource availability.
\begin{tabular}{llll}
\multicolumn{2}{l}{ Extended memonics } & \multicolumn{2}{l}{ Equivalent To } \\
\(m f f p r w z\) & RA, FRS & mfvsrwz & RA, FRS \\
\(m f v r w z\) & RA, VRS & mfvsrwz & RA, VRS +32
\end{tabular}

Special Registers Altered None

Data Layout for mfvsrwz
\(\operatorname{src}=\operatorname{VSR}[X S]\)
\begin{tabular}{|c|c|c|}
\hline unused & & unused \\
\hline
\end{tabular}
tgt \(=\) GPR[RA]
\begin{tabular}{lll}
\hline & \\
\hline 0 & 32 & 64
\end{tabular}

\section*{Move To VSR Doubleword XX1-form}
Mtvsrd
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 31 & & T & RT,RA \\
0 & 6 & & & RA & & III & \\
11 & & 179 & TX \\
31 \\
\hline
\end{tabular}
if \(T X=0\) \& MSR. FP=0 then FP Unavailablell
if \(T X=1\) \& MSR, VEC=O then Vector Unavailable()

VSR[32xTX+T].dword[0] \(\leftarrow\) GPR[RA]
VSR[32xTX+T].dword[1] \(\leftarrow\) OxUUUU_UUUU_UUUU_UUUU
Let \(X T\) be the value \(32 \times T X+T\).

The contents of GPR[RA] are placed into doubleword element 0 of VSR[ \(X T\) ].

The contents of doubleword element 1 of VSR[XT] are undefined.

For \(T X=0\), mtvsrd is treated as a Floating-Point instruction in terms of resource availability.

For \(T X=1\), mtvsrd is treated as a Vector instruction in terms of resource availability.

Extended Mnemonics

\section*{mtfprd FRT,RA}
mtvrd VRT,RA

Equivalent To
\begin{tabular}{ll} 
mtvsrd & FRT, RA \\
mtvsrd & VRT+32,RA
\end{tabular}
Special Registers Altered
None
Data Layout for mtvsrd


Move To VSR Word Algebraic XX1-form
mtvsrwa XT,RA

if \(T X=0\) \& MSR. \(F P=0\) then FP Unavailable(l)
if \(T X=1\) \& MSR. VEC \(=0\) then Vector Unavailable()
VSR[ \(32 \times T X+T]\), dwor d[0] \(\leftarrow\) EXTS64 (GPR[RA], bit [32:63])
VSR[32xTX+T].dwor d[1] \(\leftarrow\) OxUUUU_UUUU_UUUU_UUUU
Let \(X T\) be the value \(32 \times T X+T\).
The two's-complement integer in bits 32:63 of GPR[ RA] is sign-extended to 64 bits and placed into doubleword element 0 of VSR[ XT] .

The contents of doubleword element 1 of VSR[ XT] are undefined.

For TX=0, mtvsrwa is treated as a Floating-Point instruction in terms of resource availability.

For \(T X=1\), mtvsrwa is treated as a Vector instruction in terms of resource availability.
\begin{tabular}{|c|c|c|c|}
\hline \multicolumn{2}{|l|}{Extended Mnemonics} & \multicolumn{2}{|l|}{Equivalent To} \\
\hline mt fprwa & FRT, RA & mt vsrwa & FRT, RA \\
\hline mt vrwa & VRT, RA & mt vsrwa & VRT \(+32, R A\) \\
\hline
\end{tabular}

\section*{Special Registers Altered \\ None}

Data Layout for mtvsrwa


\section*{Move To VSR Word and Zero XX1-form}
Mtvsrwz
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & & T & RA \\
\hline 0 & & 6 & & 11 & \\
\hline 16 & & 21 & & 243 & TX \\
31 \\
\hline
\end{tabular}
```

if TX=0 \& MSR.FP=O then FP Unavailable||
if TX=1\& \&SR.VEC=O then Vector_Unavailable()
VSR[32xTX+T].dword[0] \& EXTZ64[GPR[RA], word[1])
VSR[32XTX+T], dwor d[1] \& OxUUUU_UUUZ_UUUU_UUUU

```

Let \(X T\) be the value \(32 \times T X+T\).
The contents of bits \(32: 63\) of GPR[RA] are placed into word element 1 of VSR[XT]. The contents of word element 0 of VSR[ XT] are set to 0 .

The contents of doubleword element 1 of VSR[ XT] are undefined.

For TX=0, mtvsrwz is treated as a Floating-Point instruction in terms of resource availability.

For \(T X=1\), \(\boldsymbol{m} t \mathbf{v s r w z}\) is treated as a Vector instruction in terms of resource availability.
\begin{tabular}{|c|c|}
\hline Extended Mnemonics & Equivalent To \\
\hline mtfprwz FRT, RA & mtvsrwz FRT, RA \\
\hline mtvrwz VRT,RA & mtvsrwz VRT+32,RA \\
\hline Special Registers Altered None & \\
\hline Data Layout for mtvsrwz & \\
\hline src = GPR[RA] & \\
\hline unused & \\
\hline \(\mathrm{tgt}=\mathrm{VSR}[\mathrm{XT}]\) & \\
\hline dword[0] & undefined \\
\hline 32 & 64 \\
\hline
\end{tabular}

\section*{Move To VSR Double Doubleword XX1-form}
mtvsrdd XT,RA,RB
\begin{tabular}{|c|cc|c|c|c|c|}
\hline 31 & 6 & T & RA & RB & & 435 \\
\hline 0 & & 6 & & & \\
\hline 12
\end{tabular}
if \(T X=0\) \& MSR. VSX \(=0\) then VSX_Unavailablell
if TX=1 \& MSR. VEC=0 then Vector_Unavailablell
VSR[32xTX+T].dword[0] \(\leftarrow(R A=0)\) ? Ox0000_0000_0000_0000: GPR[ RA]
VSR[32xTXTT]. dwor \([1] \leftarrow G P R[R B]\)
Let \(X T\) be the value \(32 \times T X+T\).
The contents of GPR[RA], or the value 0 if \(R A=0\), are placed into doubleword 0 of VSR[ XT].

The contents of GPR[RB] are placed into doubleword 1 of VSR[XT].

For \(T X=0\), mtvsrdd is treated as a Floating-Point instruction in terms of resource availability.

For \(T X=1, \boldsymbol{m} t v s r d \boldsymbol{d}\) is treated as a Vector instruction in terms of resource availability.

Special Registers Altered: None

Data Layout for mtvsrdd
\(\operatorname{src}=\) GPR[RA]
\(\square\)
SIC \(C=\) GPR[RB]
\begin{tabular}{|c|l|}
\hline & \\
\hline tgt \(=\) VSR[XT] & \\
\hline \begin{tabular}{|c|c|}
\hline. \(\operatorname{dword[0]}\) & .dword[1] \\
\hline 0 & 32
\end{tabular} \\
\hline
\end{tabular}

\section*{Move To VSR Word \& Splat XX1-form}
mtvsrws XT,RA

if \(T X=0\) \& MSR. VSX=0 then VSX Unavailable()
if \(T X=1\) \& MSR.VEC=O then Vector_Unavailable el)
VSR[32xTXTT]. word[0] \(\leftarrow\) GPR[ RA] bit [ \(32: 63]\)
VSR[32xTXTT]. word \([1] \leftarrow\) GPR[RA], bit [ \(32: 63\)
VSR[32xTX+T]. word [2] \(\leftarrow\) GPR[RA], bit [ \(32: 63]\)
VSR[32xTXTT]. word [3] \(\operatorname{GPR[~RA]~.~bit~[~} 32: 63]\)
Let \(X T\) be the value \(32 \times T X+T\).
The contents of bits \(32: 63\) of GPR[RA] are placed into each word element of VSR[ XT].

For \(T X=0\), mtvsrws is treated as a Floating-Point instruction in terms of resource availability.

For TX=1, mtvsrws is treated as a Vector instruction in terms of resource availability.

Special Registers Altered:
None

\subsection*{3.3.17 Move To/From System Register Instructions}

The Move To Condition Register Fields instruction has a preferred form; see Section 1.9.1, "Preferred Instruction Forms" on page 23. In the preferred form, the FXM field satisfies the following rule.
■ Exactly one bit of the FXM field is set to 1 .

\section*{Extended mnemonics}

Extended mnemonics are provided for the mtspr and \(m f s p r\) instructions so that they can be coded with the

SPR name as part of the mnemonic rather than as a numeric operand. An extended mnemonic is provided for the mtcrf instruction for compatibility with old software (written for a version of the architecture that precedes Version 2.00) that uses it to set the entire Condition Register. Some of these extended mnemonics are shown as examples with the relevant instructions. See Appendix C, "Assembler Extended Mnemonics" on page 795 for additional extended mnemonics.

\section*{Move To Special Purpose Register XFX-form}
mtspr SPR,RS
\begin{tabular}{|c|c|cc|cc|c|}
\hline 31 & RS & & spr & & 467 & 1 \\
0 & & 6 & 11 & & 21 & \\
31 \\
\hline
\end{tabular}
```

n}\leftarrow\mp@subsup{\operatorname{spr}}{5:9}{|}||\mp@subsup{\operatorname{spr}}{0:4}{
switch (n)
case(13): see Book III
case(808, 809, 810, 811):
default:
if length(SPR(n)) = 64 then
SPR(n) \leftarrow(RS)
else
SPR (n)}\leftarrow(RS) 32:63

```

The SPR field denotes a Special Purpose Register, encoded as shown in the table below. If the SPR field contains a value from 808 through 811, the instruction specifies a reserved SPR, and is treated as a no-op; see Section 1.3.3, "Reserved Fields, Reserved Values, and Reserved SPRs". Otherwise, unless the SPR field contains 13 (denoting the AMR), the contents of register RS are placed into the designated Special Purpose Register. For Special Purpose Registers that are 32 bits long, the low-order 32 bits of RS are placed into the SPR.

The AMR (Authority Mask Register) is used for "storage protection." This use, and operation of mtspr for the AMR, are described in Book III.
\begin{tabular}{|c|c|c|}
\hline decimal & \multicolumn{2}{c|}{\begin{tabular}{c} 
SPR \\
\(\mathbf{s p r}_{5: 9} \mathbf{s p r}_{\mathbf{0}: 4}\)
\end{tabular}} \\
\hline 1 & \begin{tabular}{c} 
Register \\
Name
\end{tabular} \\
\hline 3 & 0000000001 & XER \\
\hline 8 & 0000000011 & DSCR \\
\hline 9 & 0000001000 & LR \\
\hline 13 & 0000001001 & CTR \\
\hline 1 & 0000001101 & AMR \\
\hline
\end{tabular}

1 Note that the order of the two 5 -bit halves of the SPR number is reversed.
2 See Chapter 5 of Book II.
3 Accesses to these registers are noops; see Section 1.3.3, "Reserved Fields,
Reserved Values, and Reserved SPRs"
\begin{tabular}{|c|c|c|}
\hline decimal & \begin{tabular}{c} 
SPR \\
\(\mathbf{s p r}_{5: 9} \mathbf{s p r}_{\mathbf{0} \mathbf{4}}\)
\end{tabular} & \begin{tabular}{c} 
Register \\
Name
\end{tabular} \\
\hline 128 & 0010000000 & TFHAR \(^{2}\) \\
\hline 129 & 0010000001 & TFIAR \(^{2}\) \\
\hline 130 & 0010000010 & TEXASR \(^{2}\) \\
\hline 131 & 0010000011 & TEXASRU \(^{2}\) \\
\hline 256 & 0100000000 & VRSAVE \\
\hline 769 & 1100000001 & MMCR2 \\
\hline 770 & 1100000010 & MMCRA \\
\hline 771 & 1100000011 & PMC1 \\
\hline 772 & 1100000100 & PMC2 \\
\hline 773 & 1100000101 & PMC3 \\
\hline 774 & 1100000110 & PMC4 \\
\hline 775 & 1100000111 & PMC5 \\
\hline 776 & 1100001000 & PMC6 \\
\hline 779 & 1100001011 & MMCR0 \\
\hline 800 & 1100100000 & BESCRS \\
\hline 801 & 1100100001 & BESCRSU \\
\hline 802 & 1100100010 & BESCRR \\
\hline 803 & 1100100011 & BESCRRU \\
\hline 804 & 1100100100 & EBBHR \\
\hline 805 & 1100100101 & EBBRR \\
\hline 806 & 1100100110 & BESCR \(^{3}\) \\
\hline 808 & 1100101000 & reserved \\
\hline 809 & 1100101001 & reserved \\
\hline 8 \\
\hline 810 & 1100101010 & reserved \\
\hline 811 & 1100101011 & reserved \\
\hline
\end{tabular}

1 Note that the order of the two 5-bit halves of the SPR number is reversed.
2 See Chapter 5 of Book II.
3 Accesses to these registers are noops; see Section 1.3.3, "Reserved Fields, Reserved Values, and Reserved SPRs"
If execution of this instruction is attempted specifying an SPR number that is not shown above, one of the following occurs.
- If \(\mathrm{spr}_{0}=0\), the illegal instruction error handler is invoked.
- If \(\operatorname{spr}_{0}=1\), the system privileged instruction error handler is invoked.

If an attempt is made to execute mtspr specifying a TM SPR in other than Non-transactional state, with the exception of TFAR in suspended state, a TM Bad Thing type Program interrupt is generated.

A complete description of this instruction can be found in Book III.

\section*{Special Registers Altered:}

See above
Extended Mnemonics:
Examples of extended mnemonics for Move To Special
Purpose Register:
\begin{tabular}{llll}
\multicolumn{2}{c}{ Extended: } & \multicolumn{2}{c}{ Equivalent to: } \\
mtxer & \(R x\) & mtspr & 1,Rx \\
mtlr & \(R x\) & \(m t s p r\) & \(8, R x\) \\
mtctr & \(R x\) & \(m t s p r\) & \(9, R x\) \\
mtppr & \(R x\) & \(m t s p r\) & \(896, R x\) \\
mtppr32 & \(R x\) & \(m t s p r\) & \(898, R x\)
\end{tabular}

\section*{Programming Note}

The AMR is part of the "context" of the program (see Book III). Therefore modification of the AMR requires "synchronization" by software. For this reason, most operating systems provide a system library program that application programs can use to modify the AMR.

\section*{Compiler and Assembler Note}

For the mtspr and mfspr instructions, the SPR number coded in Assembler language does not appear directly as a 10-bit binary number in the instruction. The number coded is split into two 5-bit halves that are reversed in the instruction, with the high-order 5 bits appearing in bits 16:20 of the instruction and the low-order 5 bits in bits 11:15.

\section*{Move From Special Purpose Register XFX-form}
mfspr RT,SPR
\begin{tabular}{|c|c|cc|c|c|}
\hline 31 & RT & \multicolumn{2}{|c|}{ spr } & & 339 \\
\hline 0 & & 6 & 11 & & 21 \\
\hline 1 \\
\hline
\end{tabular}
```

n}\leftarrow\mp@subsup{\operatorname{spr}}{5:9|| |pro:4}{0
switch (n)
case(129): see Book III
case(808, 809, 810, 811):
default:
if length(SPR(n)) = 64 then
RT}\leftarrow\textrm{SPR}(\textrm{n}
else
RT}\leftarrow\mp@subsup{}{}{32}0||SPR(n

```

The SPR field denotes a Special Purpose Register, encoded as shown in the table below. If the SPR field contains 129, the instruction references the Transaction Failure Instruction Address Register (TFIAR) and the result is dependent on the privilege with which it is executed. See Book III. If the SPR field contains a value from 808 through 811, the instruction specifies a reserved SPR, and is treated as a noop; see Section 1.3.3, "Reserved Fields, Reserved Values, and Reserved SPRs". Otherwise, the contents of the designated Special Purpose Register are placed into register RT. For Special Purpose Registers that are 32 bits long, the low-order 32 bits of RT receive the contents of the Special Purpose Register and the high-order 32 bits of RT are set to zero.
\begin{tabular}{|c|c|c|}
\hline decimal & \[
\begin{gathered}
\text { SPR }^{1} \\
\text { spr }_{5: 9} \text { spr }_{0: 4}
\end{gathered}
\] & Register Name \\
\hline 1 & 0000000001 & XER \\
\hline 3 & 0000000011 & DSCR \\
\hline 8 & 0000001000 & LR \\
\hline 9 & 0000001001 & CTR \\
\hline 13 & 0000001101 & AMR \\
\hline 128 & 0010000000 & TFHAR \({ }^{4}\) \\
\hline 129 & 0010000001 & TFIAR \({ }^{4}\) \\
\hline 130 & 0010000010 & TEXASR \({ }^{4}\) \\
\hline 131 & 0010000011 & TEXASRU4 \\
\hline 136 & 0010001000 & CTRL \\
\hline 256 & 0100000000 & VRSAVE \\
\hline 259 & 0100000011 & SPRG3 \\
\hline 268 & 0100001100 & TB \({ }^{2}\) \\
\hline 269 & 0100001101 & TBU \({ }^{2}\) \\
\hline 768 & 1100000000 & SIER \\
\hline 769 & 1100000001 & MMCR2 \\
\hline 770 & 1100000010 & MMCRA \\
\hline 771 & 1100000011 & PMC1 \\
\hline \multicolumn{3}{|l|}{\multirow[t]{4}{*}{\begin{tabular}{l}
Note that the order of the two 5-bit halves of the SPR number is reversed. \\
2 See Chapter 6 of Book II \\
3 Accesses to these SPRs are noops; see Section 1.3.3, "Reserved Fields, Reserved Values, and Reserved SPRs". \\
4 See Chapter 5 of Book II.
\end{tabular}}} \\
\hline & & \\
\hline & & \\
\hline & & \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|}
\hline decimal & \[
\begin{gathered}
\text { SPR }^{1} \\
\text { spr }_{5: 9} \text { spr }_{0: 4}
\end{gathered}
\] & Register Name \\
\hline 772 & 1100000100 & PMC2 \\
\hline 773 & 1100000101 & PMC3 \\
\hline 774 & 1100000110 & PMC4 \\
\hline 775 & 1100000111 & PMC5 \\
\hline 776 & 1100001000 & PMC6 \\
\hline 779 & 1100001011 & MMCR0 \\
\hline 780 & 1100001100 & SIAR \\
\hline 781 & 1100001101 & SDAR \\
\hline 782 & 1100001110 & MMCR1 \\
\hline 800 & 1100100000 & BESCRS \\
\hline 801 & 1100100001 & BESCRSU \\
\hline 802 & 1100100010 & BESCRR \\
\hline 803 & 1100100011 & BESCRRU \\
\hline 804 & 1100100100 & EBBHR \\
\hline 805 & 1100100101 & EBBRR \\
\hline 806 & 1100100110 & BESCR \\
\hline 808 & 1100101000 & reserved \({ }^{3}\) \\
\hline 809 & 1100101001 & reserved \({ }^{3}\) \\
\hline 810 & 1100101010 & reserved \({ }^{3}\) \\
\hline 811 & 1100101011 & reserved \({ }^{3}\) \\
\hline 813 & 1100101101 & LMRR \\
\hline 814 & 1100101110 & LMSER \\
\hline 815 & 1100101111 & TAR \\
\hline 896 & 1110000000 & PPR \({ }^{10}\) \\
\hline 898 & 1110000010 & PPR32 \\
\hline \multicolumn{3}{|l|}{1 Note that the order of the two 5-bit halves of the SPR number is reversed.} \\
\hline \multicolumn{3}{|l|}{2 See Chapter 6 of Book II} \\
\hline \multicolumn{3}{|l|}{\multirow[t]{4}{*}{\begin{tabular}{|ll}
3 & \(\begin{array}{l}\text { Accesses to these SPRs are noops; see } \\
\text { Section 1.3.3, "Reserved Fields, Reserved } \\
\text { Values, and Reserved SPRs". }\end{array}\) \\
\(4 \begin{array}{l}\text { See Chapter } 5 \text { of Book II. }\end{array}\)
\end{tabular}}} \\
\hline & & \\
\hline & & \\
\hline & & \\
\hline
\end{tabular}

If execution of this instruction is attempted specifying an SPR number that is not shown above, one of the following occurs.
- If \(\mathrm{spr}_{0}=0\), the illegal instruction error handler is invoked.
- If \(\mathrm{spr}_{0}=1\), the system privileged instruction error handler is invoked.

A complete description of this instruction can be found in Book III.

\section*{Special Registers Altered:}

None

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Move From Special Purpose Register:
\begin{tabular}{ll}
\multicolumn{2}{c}{ Extended: } \\
mfxer & \(R x\) \\
\(m f l r\) & \(R x\) \\
\(m f c t r\) & \(R x\)
\end{tabular}

Equivalent to:
mfspr Rx, 1
mfspr Rx,8
mfspr Rx,9

\footnotetext{
- Note

See the Notes that appear with mtspr.
}

\section*{Move to CR from XER Extended X-form}
morxrx BF
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 31 & BF & II & III & & III & & 576 \\
\hline 0 & & 6 & 9 & 11 & 16 & & 21 \\
31
\end{tabular}
\(\mathrm{CR}_{4 \times \mathrm{BF}+32: 4 \times \mathrm{BF}+35} \leftarrow \mathrm{XER}_{\mathrm{OV}}\) ov32 CA CA32
The contents of the OV, OV32, CA, and CA32 are copied to Condition Register field BF.

Special Registers Altered:
CR field BF

\section*{Move To One Condition Register Field} XFX-form
mtocrf FXM,RS

```

count \leftarrow0
do i = 0 to 7
if FXM i
n}\leftarrow
count }\leftarrow\mathrm{ count + 1
if count = 1 then
CR
else CR }\leftarrow\mathrm{ undefined

```

If exactly one bit of the FXM field is set to 1 , let n be the position of that bit in the field ( \(0 \leq n \leq 7\) ). The contents of bits \(4 \times n+32: 4 \times n+35\) of register RS are placed into CR field \(n\) (CR bits \(4 \times n+32: 4 \times n+35\) ). Otherwise, the contents of the Condition Register are undefined.

\section*{Special Registers Altered:}

CR field selected by FXM

\section*{Move To Condition Register Fields} XFX-form
mtcrf FXM,RS
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & 0 & FXM & \(/\) & & 144 \\
\hline 0 & & 6 & 11 & 12 & 20 & 21 \\
& & & \\
\hline
\end{tabular}
```

mask \leftarrow \& }\mp@subsup{}{(}{(\mp@subsup{FXXM}{0}{})| | | (FXM (
CR}\leftarrow((RS) 32:63\& mask) | (CR\& ᄀmask)

```

The contents of bits 32:63 of register RS are placed into the Condition Register under control of the field mask specified by FXM. The field mask identifies the 4 -bit fields affected. Let i be an integer in the range 0-7. If \(\mathrm{FXM}_{\mathrm{i}}=1\) then CR field i (CR bits \(4 \times i+32: 4 \times i+35\) ) is set to the contents of the corresponding field of the low-order 32 bits of RS.

\section*{Special Registers Altered:}

CR fields selected by mask

\section*{Extended Mnemonics:}

Example of extended mnemonics for Move To Condition Register Fields:
\begin{tabular}{lll} 
Extended: & Equivalent to: \\
mtcr & \(R x\) & mtcrf \\
& \(0 x F F, R x\)
\end{tabular}

\section*{Move From One Condition Register Field XFX-form}
```

mfocrf
RT,FXM

| 31 | RT | 1 | FXM | $/$ |  | 19 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 12 | 20 | 21 |  |

```
```

RT \leftarrow undefined

```
RT \leftarrow undefined
count }\leftarrow
count }\leftarrow
do i = 0 to 7
do i = 0 to 7
    if FXM 
    if FXM 
        n}\leftarrow\textrm{i
        n}\leftarrow\textrm{i
        count \leftarrow count + 1
        count \leftarrow count + 1
if count = 1 then
if count = 1 then
    RT}\leftarrow\mp@subsup{}{}{64}
    RT}\leftarrow\mp@subsup{}{}{64}
    RT
```

    RT
    ```

If exactly one bit of the FXM field is set to 1 , let \(n\) be the position of that bit in the field ( \(0 \leq \mathrm{n} \leq 7\) ). The contents of CR field \(n(C R\) bits \(4 \times n+32: 4 \times n+35\) ) are placed into bits \(4 \times n+32: 4 \times n+35\) of register RT, and the contents of the remaining bits of register RT are undefined. Otherwise, the contents of register RT are undefined.

If exactly one bit of the FXM field is set to 1, the contents of the remaining bits of register RT are set to 0 's instead of being undefined as specified above.

\section*{Special Registers Altered:}

None

\section*{Programming Note}

Warning: mfocrf is not backward compatible with processors that comply with versions of the architecture that precede Version 3.0. Such processors may not set to 0 the bits of register RT that do not correspond to the specified CR field. If programs that depend on this clearing behavior are run on such processors, the programs may get incorrect results.

The POWER4, POWER5, POWER7 and POWER8 processors set to O's all bytes of register RT other than the byte that contains the specified CR field. In the byte that contains the CR field, bits other than those containing the CR field may or may not be set to 0 s .

Move From Condition Register XFX-form
RT
mfcr
\begin{tabular}{|c|c|c|c|cc|c|}
\hline 31 & RT & 0 & & I/I & & 19 \\
\hline 0 & & 6 & & 11 & 12 & \\
\hline
\end{tabular}
\(\mathrm{RT} \leftarrow{ }^{32} 0 \| \mathrm{CR}\)
The contents of the Condition Register are placed into \(R T_{32: 63} . \mathrm{RT}_{0: 31}\) are set to 0 .

\section*{Special Registers Altered:}

None

\section*{Set Boolean X-form \\ setb RT,BFA}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \({ }_{0} 31\) & 6 & RT & \({ }_{11}\) & & 16 & III & 21 & 128 & \(1 \begin{aligned} & 1 \\ & 31\end{aligned}\) \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { if } C R_{4 \times 8 F A+32}=1 \text { then } \\
& \text { RT } \leftarrow \text { OXFFFF_FFF_FFFFFFFF } \\
& \text { else if } \mathrm{CR}_{4 \times 8 \mathrm{FA}+33}=1 \text { then } \\
& R T \leftarrow 0 \times 0000000000000001 \\
& \text { else } \\
& R T \leftarrow 0 \times 0000 \_0000 \_000000000
\end{aligned}
\]

If the contents of bit 0 of \(C R\) field \(B F A\) are equal to \(0 b 1\), the contents of register RT are set to OXFFFF_FFF_FFF_FFFF.

Otherwise, if the contents of bit 1 of \(C R\) field \(B F A\) are equal to Ob1, the contents of register RT are set to \(0 \times 0000 \_0000 \_0000 \_0001\).

Otherwise, the contents of register RT are set to \(0 \times 0000 \_0000 \_0000 \_0000\).

\section*{Special Registers Altered:}

None

Version 3.0

\title{
Chapter 4. Floating-Point Facility
}

\subsection*{4.1 Floating-Point Facility Overview}

This chapter describes the registers and instructions that make up the Floating-Point Facility.
The processor (augmented by appropriate software support, where required) implements a floating-point system compliant with the ANSI/IEEE Standard 754-1985, "IEEE Standard for Binary Floating-Point Arithmetic" (hereafter referred to as "the IEEE standard"). That standard defines certain required "operations" (addition, subtraction, etc.). Herein, the term "floating-point operation" is used to refer to one of these required operations and to additional operations defined (e.g., those performed by Multiply-Add or Reciprocal Estimate instructions). A Non-IEEE mode is also provided. This mode, which may produce results not in strict compliance with the IEEE standard, allows shorter latency.
Instructions are provided to perform arithmetic, rounding, conversion, comparison, and other operations in floating-point registers; to move floating-point data between storage and these registers; and to manipulate the Floating-Point Status and Control Register explicitly.

These instructions are divided into two categories.

\section*{- computational instructions}

The computational instructions are those that perform addition, subtraction, multiplication, division, extracting the square root, rounding, conversion, comparison, and combinations of these operations. These instructions provide the floating-point operations. They place status information into the Floating-Point Status and Control Register. They are the instructions described in Sections 4.6.6 through 4.6.8.
- non-computational instructions

The non-computational instructions are those that perform loads and stores, move the contents of a floating-point register to another floating-point register possibly altering the sign, manipulate the Floating-Point Status and Control Register explic-
itly, and select the value from one of two float-ing-point registers based on the value in a third floating-point register. The operations performed by these instructions are not considered float-ing-point operations. With the exception of the instructions that manipulate the Floating-Point Status and Control Register explicitly, they do not alter the Floating-Point Status and Control Register. They are the instructions described in Sections 4.6.2 through 4.6.5, and 4.6.10.

A floating-point number consists of a signed exponent and a signed significand. The quantity expressed by this number is the product of the significand and the number \(2^{\text {exponent }}\). Encodings are provided in the data format to represent finite numeric values, \(\pm\) Infinity, and values that are "Not a Number" (NaN). Operations involving infinities produce results obeying traditional mathematical conventions. NaNs have no mathematical interpretation. Their encoding permits a variable diagnostic information field. They may be used to indicate such things as uninitialized variables and can be produced by certain invalid operations.

There is one class of exceptional events that occur during instruction execution that is unique to the Float-ing-Point Facility: the Floating-Point Exception. Floating-point exceptions are signaled with bits set in the Floating-Point Status and Control Register (FPSCR). They can cause the system floating-point enabled exception error handler to be invoked, precisely or imprecisely, if the proper control bits are set.

\section*{Floating-Point Exceptions}

The following floating-point exceptions are detected by the processor:
```

■ Invalid Operation Exception
SNaN
Infinity-Infinity
Infinity\divInfinity (VXIDI)
Zero\divZero
Infinity\timesZero
Invalid Compare
Software-Defined Condition (VXSOFT)
Invalid Square Root

Invalid Integer Convert<br>- Zero Divide Exception<br>- Overflow Exception<br>- Underflow Exception<br>- Inexact Exception (VXCVI)

Each floating-point exception, and each category of Invalid Operation Exception, has an exception bit in the FPSCR. In addition, each floating-point exception has a corresponding enable bit in the FPSCR. See Section 4.2.2, "Floating-Point Status and Control Register" on page 124 for a description of these exception and enable bits, and Section 4.4, "Floating-Point Exceptions" on page 132 for a detailed discussion of floating-point exceptions, including the effects of the enable bits.

### 4.2 Floating-Point Facility Registers

### 4.2.1 Floating-Point Registers

Implementations of this architecture provide 32 float-ing-point registers (FPRs). The floating-point instruction formats provide 5 -bit fields for specifying the FPRs to be used in the execution of the instruction. The FPRs are numbered 0-31. See Figure 48 on page 124.
Each FPR contains 64 bits that support the float-ing-point double format. Every instruction that interprets the contents of an FPR as a floating-point value uses the floating-point double format for this interpretation.
The computational instructions, and the Move and Select instructions, operate on data located in FPRs and, with the exception of the Compare instructions, place the result value into an FPR and optionally (when Rc=1) place status information into the Condition Register.

Load Double and Store Double instructions are provided that transfer 64 bits of data between storage and the FPRs with no conversion. Load Single instructions are provided to transfer and convert floating-point values in floating-point single format from storage to the same value in floating-point double format in the FPRs. Store Single instructions are provided to transfer and convert floating-point values in floating-point double format from the FPRs to the same value in float-ing-point single format in storage.
Instructions are provided that manipulate the Float-ing-Point Status and Control Register and the Condition Register explicitly. Some of these instructions copy data from an FPR to the Floating-Point Status and Control Register or vice versa.

The computational instructions and the Select instruction accept values from the FPRs in double format. For single-precision arithmetic instructions, all input values must be representable in single format; if they are not,
the result placed into the target FPR, and the setting of status bits in the FPSCR and in the Condition Register (if Rc=1), are undefined.

| FPR 0 |  |
| :---: | :---: |
| FPR 1 |  |
| $\cdots$ |  |
| $\cdots$ |  |
| FPR 30 |  |
| FPR 31 |  |
| 0 |  |

Figure 48. Floating-Point Registers

### 4.2.2 Floating-Point Status and Control Register

The Floating-Point Status and Control Register (FPSCR) controls the handling of floating-point exceptions and records status resulting from the float-ing-point operations. Bits $32: 55$ are status bits. Bits 56:63 are control bits.

The exception bits in the FPSCR (bits $35: 44,53: 55$ ) are sticky; that is, once set to 1 they remain set to 1 until they are set to 0 by an mcrfs, mtfsfi, mtfsf, or mtfsb0 instruction. The exception summary bits in the FPSCR (FX, FEX, and VX, which are bits 32:34) are not considered to be "exception bits", and only FX is sticky.
FEX and VX are simply the ORs of other FPSCR bits. Therefore these two bits are not listed among the FPSCR bits affected by the various instructions.


## Figure 49. Floating-Point Status and Control Register

The bit definitions for the FPSCR are as follows.

## Bit(s) Description

0:31 Reserved
$32 \quad$ Floating-Point Exception Summary (FX)
Every floating-point instruction, except mtfsfi and $\boldsymbol{m t f s f}$, implicitly sets FPSCR ${ }_{F X}$ to 1 if that instruction causes any of the floating-point exception bits in the FPSCR to change from 0 to 1. merfs, mtfsfi, mtfsf, mtfsbO, and mtfsb1 can alter FPSCR $_{F X}$ explicitly.

## Programming Note

FPSCR $_{\text {FX }}$ is defined not to be altered implicitly by mtfsfi and mtfsf because permitting these instructions to alter $\mathrm{FPSCR}_{\text {FX }}$ implicitly could cause a paradox. An example is an mtfsfi or mtfsf instruction that supplies 0 for FPSCR $_{\text {FX }}$ and 1 for $\mathrm{FPSCR}_{\mathrm{OX}}$, and is executed when $\mathrm{FPSCR}_{\mathrm{Ox}}=0$. See also the Programming Notes with the definition of these two instructions.

Floating-Point Enabled Exception Summary (FEX)
This bit is the OR of all the floating-point exception bits masked by their respective enable bits. mcrfs, mtfsfi, mtfsf, mtfsb0, and mtfsb1 cannot alter FPSCR FEX explicitly.
Floating-Point Invalid Operation Exception Summary (VX)
This bit is the OR of all the Invalid Operation exception bits. mcrfs, mtfsfi, mtfsf, mtfsb0, and mtfsb1 cannot alter FPSCR ${ }_{\mathrm{Vx}}$ explicitly.
Floating-Point Overflow Exception (OX) See Section 4.4.3, "Overflow Exception" on page 135.
Floating-Point Underflow Exception (UX) See Section 4.4.4, "Underflow Exception" on page 136.
Floating-Point Zero Divide Exception (ZX)
See Section 4.4.2, "Zero Divide Exception" on page 134.

Floating-Point Inexact Exception (XX)
See Section 4.4.5, "Inexact Exception" on page 136.

FPSCR $_{X X}$ is a sticky version of FPSCR $_{\text {FI }}$ (see below). Thus the following rules completely describe how FPSCR $_{\mathrm{Xx}}$ is set by a given instruction.

■ If the instruction affects FPSCR $_{\text {FI }}$, the new value of FPSCR $_{x x}$ is obtained by ORing the old value of FPSCR $X_{X X}$ with the new value of $\mathrm{FPSCR}_{\mathrm{FI}}$.
■ If the instruction does not affect FPSCR $_{\text {Fl }}$, the value of FPSCR $_{X X}$ is unchanged.
Floating-Point Invalid Operation Exception (SNaN) (VXSNAN)
See Section 4.4.1, "Invalid Operation Exception" on page 134.

Floating-Point Invalid Operation Exception ( $\infty$ - $\infty$ ) (VXISI)
See Section 4.4.1.

Floating-Point Invalid Operation Exception $(\infty \div \infty)$ (VXIDI) See Section 4.4.1.
Floating-Point Invalid Operation Exception $(0 \div 0)(V X Z D Z)$
See Section 4.4.1.
$3 \quad$ Floating-Point Invalid Operation Exception ( $\infty \times 0$ ) (VXIMZ)
See Section 4.4.1.
Floating-Point Invalid Operation Exception (Invalid Compare) (VXVC)
See Section 4.4.1.
Floating-Point Fraction Rounded (FR)
The last Arithmetic or Rounding and Conversion instruction incremented the fraction during rounding. See Section 4.3.6, "Rounding" on page 131. This bit is not sticky.
Floating-Point Fraction Inexact (FI)
The last Arithmetic or Rounding and Conversion instruction either produced an inexact result during rounding or caused a disabled Overflow Exception. See Section 4.3.6. This bit is not sticky.

See the definition of $\mathrm{FPSCR}_{X X}$, above, regarding the relationship between $\mathrm{FPSCR}_{\text {FI }}$ and FPSCR ${ }_{X X}$.

## Floating-Point Result Flags (FPRF)

Arithmetic, rounding, and Convert From Integer instructions set this field based on the result placed into the target register and on the target precision, except that if any portion of the result is undefined then the value placed into FPRF is undefined. Floating-point Compare instructions set this field based on the relative values of the operands being compared. For Convert To Integer instructions, the value placed into FPRF is undefined. Additional details are given below.

## Programming Note

A single-precision operation that produces a denormalized result sets FPRF to indicate a denormalized number. When possible, single-precision denormalized numbers are represented in normalized double format in the target register.

Floating-Point Result Class Descriptor (C) Arithmetic, rounding, and Convert From Integer instructions may set this bit with the FPCC bits, to indicate the class of the result as shown in Figure 50 on page 127.

Floating-Point Condition Code (FPCC)

Floating-point Compare instructions set one of
$48 \quad$ Floating-Point Less Than or Negative (FL or <)

49 Floating-Point Greater Than or Positive (FG or >)

50 Floating-Point Equal or Zero (FE or =)
$51 \quad$ Floating-Point Unordered or NaN (FU or ?)
52 Reserved
$53 \quad$ Floating-Point Invalid Operation Exception (Software-Defined Condition)
(VXSOFT)
This bit can be altered only by mcrfs, mtfsfi, $\boldsymbol{m t f s f}, \boldsymbol{m t f s b 0}$, or mtfsb1. See Section 4.4.1.

## Programming Note

FPSCR $_{\text {VXSOFT }}$ can be used by software to indicate the occurrence of an arbitrary, software-defined, condition that is to be treated as an Invalid Operation Exception. For example, the bit could be set by a program that computes a base 10 logarithm if the supplied input is negative.

54 Floating-Point Invalid Operation Exception (Invalid Square Root) (VXSQRT) See Section 4.4.1.

55 Floating-Point Invalid Operation Exception (Invalid Integer Convert) (VXCVI) See Section 4.4.1.

56 Floating-Point Invalid Operation Exception Enable (VE)
See Section 4.4.1.
57 Floating-Point Overflow Exception Enable (OE)
See Section 4.4.3, "Overflow Exception" on page 135.

58 Floating-Point Underflow Exception Enable (UE)
See Section 4.4.4, "Underflow Exception" on page 136.
59 Floating-Point Zero Divide Exception Enable (ZE)
See Section 4.4.2, "Zero Divide Exception" on page 134.
60 Floating-Point Inexact Exception Enable (XE)

See Section 4.4.5, "Inexact Exception" on page 136.
61 Floating-Point Non-IEEE Mode (NI)
Floating-point non-IEEE mode is optional. If floating-point non-IEEE mode is not implemented, this bit is treated as reserved, and the remainder of the definition of this bit does not apply.
If floating-point non-IEEE mode is implemented, this bit has the following meaning.
0 The processor is not in floating-point non-IEEE mode (i.e., all floating-point operations conform to the IEEE standard).
1 The processor is in floating-point non-IEEE mode.

When the processor is in floating-point non-IEEE mode, the remaining FPSCR bits may have meanings different from those given in this document, and floating-point operations need not conform to the IEEE standard. The effects of executing a given floating-point instruction with $\mathrm{FPSCR}_{\mathrm{NI}}=1$, and any additional requirements for using non-IEEE mode, are implementation-dependent. The results of executing a given instruction in non-IEEE mode may vary between implementations, and between different executions on the same implementation.

## Programming Note

When the processor is in floating-point non-IEEE mode, the results of float-ing-point operations may be approximate, and performance for these operations may be better, more predictable, or less data-dependent than when the processor is not in non-IEEE mode. For example, in non-IEEE mode an implementation may return 0 instead of a denormalized number, and may return a large number instead of an infinity.

Floating-Point Rounding Control (RN) See Section 4.3.6, "Rounding" on page 131.
00 Round to Nearest
01 Round toward Zero
10 Round toward +Infinity
11 Round toward -Infinity

| Result Flags | Result Value Class |
| :---: | :---: |
| $\mathrm{C}<>=$ ? |  |
| 10001 | Quiet NaN |
| 01001 | - Infinity |
| 01000 | - Normalized Number |
| 11000 | - Denormalized Number |
| 10010 | - Zero |
| 00010 | + Zero |
| 10100 | + Denormalized Number |
| 00100 | + Normalized Number |
| 00101 | + Infinity |

Figure 50. Floating-Point Result Flags

### 4.3 Floating-Point Data

### 4.3.1 Data Format

This architecture defines the representation of a float-ing-point value in two different binary fixed-length formats. The format may be a 32-bit single format for a single-precision value or a 64-bit double format for a double-precision value. The single format may be used for data in storage. The double format may be used for data in storage and for data in floating-point registers.
The lengths of the exponent and the fraction fields differ between these two formats. The structure of the single and double formats is shown below.


Figure 51. Floating-point single format


Figure 52. Floating-point double format
Values in floating-point format are composed of three fields:

```
S sign bit
EXP exponent+bias
FRACTION fraction
```

Representation of numeric values in the floating-point formats consists of a sign bit (S), a biased exponent (EXP), and the fraction portion (FRACTION) of the significand. The significand consists of a leading implied bit concatenated on the right with the FRACTION. This leading implied bit is 1 for normalized numbers and 0 for denormalized numbers and is located in the unit bit position (i.e., the first bit to the left of the binary point). Values representable within the two floating-point for-
mats can be specified by the parameters listed in Figure 53.

|  | Format |  |
| :--- | :---: | :---: |
|  | Single | Double |
|  |  |  |
| Exponent Bias | +127 | +1023 |
| Maximum Exponent | +127 | +1023 |
| Minimum Exponent | -126 | -1022 |
|  |  |  |
| Widths (bits) | 32 | 64 |
| $\quad$ Format | 1 | 1 |
| Sign | 8 | 11 |
| Exponent | 23 | 52 |
| Fraction | 24 | 53 |
| Significand |  |  |
|  |  |  |

Figure 53. IEEE floating-point fields
The architecture requires that the FPRs of the Float-ing-Point Facility support the floating-point double format only.

### 4.3.2 Value Representation

This architecture defines numeric and non-numeric values representable within each of the two supported formats. The numeric values are approximations to the real numbers and include the normalized numbers, denormalized numbers, and zero values. The non-numeric values representable are the infinities and the Not a Numbers (NaNs). The infinities are adjoined to the real numbers, but are not numbers themselves, and the standard rules of arithmetic do not hold when they are used in an operation. They are related to the real numbers by order alone. It is possible however to define restricted operations among numbers and infinities as defined below. The relative location on the real number line for each of the defined entities is shown in Figure 54.


Figure 54. Approximation to real numbers
The NaNs are not related to the numeric values or infinities by order or value but are encodings used to convey diagnostic information such as the representation of uninitialized variables.
The following is a description of the different float-ing-point values defined in the architecture:

## Binary floating-point numbers

Machine representable values used as approximations to real numbers. Three categories of numbers are supported: normalized numbers, denormalized numbers, and zero values.

Normalized numbers ( $\pm$ NOR)
These are values that have a biased exponent value in the range:

1 to 254 in single format
1 to 2046 in double format
They are values in which the implied unit bit is 1 . Normalized numbers are interpreted as follows:

$$
\mathrm{NOR}=(-1)^{\mathrm{s}} \times 2^{\mathrm{E}} \times \text { (1.fraction) }
$$

where $s$ is the sign, $E$ is the unbiased exponent, and 1.fraction is the significand, which is composed of a leading unit bit (implied bit) and a fraction part.
The ranges covered by the magnitude ( $M$ ) of a normalized floating-point number are approximately equal to:

## Single Format:

$$
1.2 \times 10^{-38} \leq \mathrm{M} \leq 3.4 \times 10^{38}
$$

## Double Format:

$$
2.2 \times 10^{-308} \leq \mathrm{M} \leq 1.8 \times 10^{308}
$$

## Zero values ( $\pm 0$ )

These are values that have a biased exponent value of zero and a fraction value of zero. Zeros can have a positive or negative sign. The sign of zero is ignored by comparison operations (i.e., comparison regards +0 as equal to -0 ).

## Denormalized numbers ( $\pm$ DEN)

These are values that have a biased exponent value of zero and a nonzero fraction value. They are nonzero numbers smaller in magnitude than the representable normalized numbers. They are values in which the implied unit bit is 0 . Denormalized numbers are interpreted as follows:

$$
\text { DEN }=(-1)^{\mathrm{S}} \times 2^{\mathrm{Emin}} \times \text { (0.fraction) }
$$

where Emin is the minimum representable exponent value (-126 for single-precision, -1022 for double-precision).

Infinities ( $\pm \infty$ )
These are values that have the maximum biased exponent value:

255 in single format
2047 in double format
and a zero fraction value. They are used to approximate values greater in magnitude than the maximum normalized value.
Infinity arithmetic is defined as the limiting case of real arithmetic, with restricted operations defined among numbers and infinities. Infinities and the real numbers can be related by ordering in the affine sense:

$$
-\infty<\text { every finite number }<+\infty
$$

Arithmetic on infinities is always exact and does not signal any exception, except when an exception occurs
due to the invalid operations as described in Section 4.4.1, "Invalid Operation Exception" on page 134.
For comparison operations, +Infinity compares equal to +Infinity and -Infinity compares equal to -Infinity.
Not a Numbers (NaNs)
These are values that have the maximum biased exponent value and a nonzero fraction value. The sign bit is ignored (i.e., NaNs are neither positive nor negative). If the high-order bit of the fraction field is 0 then the NaN is a Signaling NaN ; otherwise it is a Quiet NaN .

Signaling NaNs are used to signal exceptions when they appear as operands of computational instructions.
Quiet NaNs are used to represent the results of certain invalid operations, such as invalid arithmetic operations on infinities or on NaNs, when Invalid Operation Exception is disabled (FPSCR ${ }_{V E}=0$ ). Quiet NaNs propagate through all floating-point operations except ordered comparison, Floating Round to Single-Precision, and conversion to integer. Quiet NaNs do not signal exceptions, except for ordered comparison and conversion to integer operations. Specific encodings in QNaNs can thus be preserved through a sequence of floating-point operations, and used to convey diagnostic information to help identify results from invalid operations.

When a QNaN is the result of a floating-point operation because one of the operands is a NaN or because a QNaN was generated due to a disabled Invalid Operation Exception, then the following rule is applied to determine the NaN with the high-order fraction bit set to 1 that is to be stored as the result.

```
if (FRA) is a NaN
    then FRT \leftarrow(FRA)
    else if (FRB) is a NaN
        then if instruction is frsp
            then FRT}\leftarrow(FRB\mp@subsup{)}{0:34 II 290}{
            else FRT \leftarrow(FRB)
        else if (FRC) is a NaN
            then FRT}\leftarrow(FRC
            else if generated QNaN
                then FRT}\leftarrow\mathrm{ generated QNaN
```

If the operand specified by FRA is a NaN, then that NaN is stored as the result. Otherwise, if the operand specified by FRB is a NaN (if the instruction specifies an FRB operand), then that NaN is stored as the result, with the low-order 29 bits of the result set to 0 if the instruction is frsp. Otherwise, if the operand specified by FRC is a NaN (if the instruction specifies an FRC operand), then that NaN is stored as the result. Otherwise, if a QNaN was generated due to a disabled Invalid Operation Exception, then that QNaN is stored as the result. If a QNaN is to be generated as a result, then the QNaN generated has a sign bit of 0 , an exponent field of all 1 s , and a high-order fraction bit of 1 with all other fraction bits 0 . Any instruction that generates a QNaN as the result of a disabled Invalid Operation

Exception generates this QNaN (i.e., 0x7FF8_0000_0000_0000).
A double-precision NaN is considered to be representable in single format if and only if the low-order 29 bits of the double-precision NaN's fraction are zero.

### 4.3.3 Sign of Result

The following rules govern the sign of the result of an arithmetic, rounding, or conversion operation, when the operation does not yield an exception. They apply even when the operands or results are zeros or infinities.

- The sign of the result of an add operation is the sign of the operand having the larger absolute value. If both operands have the same sign, the sign of the result of an add operation is the same as the sign of the operands. The sign of the result of the subtract operation $x-y$ is the same as the sign of the result of the add operation $\mathrm{x}+(-\mathrm{y})$.
When the sum of two operands with opposite sign, or the difference of two operands with the same sign, is exactly zero, the sign of the result is positive in all rounding modes except Round toward - Infinity, in which mode the sign is negative.
- The sign of the result of a multiply or divide operation is the Exclusive OR of the signs of the operands.
- The sign of the result of a Square Root or Reciprocal Square Root Estimate operation is always positive, except that the square root of -0 is -0 and the reciprocal square root of -0 is - Infinity.
- The sign of the result of a Round to Single-Precision, or Convert From Integer, or Round to Integer operation is the sign of the operand being converted.

For the Multiply-Add instructions, the rules given above are applied first to the multiply operation and then to the add or subtract operation (one of the inputs to the add or subtract operation is the result of the multiply operation).

### 4.3.4 Normalization and Denormalization

The intermediate result of an arithmetic or frsp instruction may require normalization and/or denormalization as described below. Normalization and denormalization do not affect the sign of the result.

When an arithmetic or rounding instruction produces an intermediate result which carries out of the significand, or in which the significand is nonzero but has a leading zero bit, it is not a normalized number and must be normalized before it is stored. For the carry-out case, the significand is shifted right one bit, with a one shifted into the leading significand bit, and the exponent is incre-
mented by one. For the leading-zero case, the significand is shifted left while decrementing its exponent by one for each bit shifted, until the leading significand bit becomes one. The Guard bit and the Round bit (see Section 4.5.1, "Execution Model for IEEE Operations" on page 137) participate in the shift with zeros shifted into the Round bit. The exponent is regarded as if its range were unlimited.

After normalization, or if normalization was not required, the intermediate result may have a nonzero significand and an exponent value that is less than the minimum value that can be represented in the format specified for the result. In this case, the intermediate result is said to be "Tiny" and the stored result is determined by the rules described in Section 4.4.4, "Underflow Exception". These rules may require denormalization.

A number is denormalized by shifting its significand right while incrementing its exponent by 1 for each bit shifted, until the exponent is equal to the format's minimum value. If any significant bits are lost in this shifting process then "Loss of Accuracy" has occurred (See Section 4.4.4, "Underflow Exception" on page 136) and Underflow Exception is signaled.

### 4.3.5 Data Handling and Precision

Most of the Floating-Point Facility Architecture, including all computational, Move, and Select instructions, use the floating-point double format to represent data in the FPRs. Single-precision and integer-valued operands may be manipulated using double-precision operations. Instructions are provided to coerce these values from a double format operand. Instructions are also provided for manipulations which do not require dou-ble-precision. In addition, instructions are provided to access a true single-precision representation in storage, and a fixed-point integer representation in GPRs.

### 4.3.5.1 Single-Precision Operands

For single format data, a format conversion from single to double is performed when loading from storage into an FPR and a format conversion from double to single is performed when storing from an FPR to storage. No floating-point exceptions are caused by these instructions. An instruction is provided to explicitly convert a double format operand in an FPR to single-precision. Floating-point single-precision is enabled with four types of instruction.

## 1. Load Floating-Point Single

This form of instruction accesses a single-precision operand in single format in storage, converts it to double format, and loads it into an FPR. No floating-point exceptions are caused by these instructions.

## 2. Round to Floating-Point Single-Precision

The Floating Round to Single-Precision instruction rounds a double-precision operand to single-precision, checking the exponent for single-precision range and handling any exceptions according to respective enable bits, and places that operand into an FPR in double format. For results produced by single-precision arithmetic instructions, sin-gle-precision loads, and other instances of the Floating Round to Single-Precision instruction, this operation does not alter the value.
3. Single-Precision Arithmetic Instructions

This form of instruction takes operands from the FPRs in double format, performs the operation as if it produced an intermediate result having infinite precision and unbounded exponent range, and then coerces this intermediate result to fit in single format. Status bits, in the FPSCR and optionally in the Condition Register, are set to reflect the sin-gle-precision result. The result is then converted to double format and placed into an FPR. The result lies in the range supported by the single format.

If any input value is not representable in single format and either $\mathrm{OE}=1$ or $\mathrm{UE}=1$, the result placed into the target FPR, and the setting of status bits in the FPSCR and in the Condition Register (if $\mathrm{Rc}=1$ ), are undefined.

For fres[.] or frsqrtes[.], if the input value is finite and has an unbiased exponent greater than +127 , the input value is interpreted as an Infinity.

## 4. Store Floating-Point Single

This form of instruction converts a double-precision operand to single format and stores that operand into storage. No floating-point exceptions are caused by these instructions. (The value being stored is effectively assumed to be the result of an instruction of one of the preceding three types.)

When the result of a Load Floating-Point Single, Floating Round to Single-Precision, or single-precision arithmetic instruction is stored in an FPR, the low-order 29 FRACTION bits are zero.

## Programming Note

The Floating Round to Single-Precision instruction is provided to allow value conversion from dou-ble-precision to single-precision with appropriate exception checking and rounding. This instruction should be used to convert double-precision float-ing-point values (produced by double-precision load and arithmetic instructions and by fcfid) to sin-gle-precision values prior to storing them into single format storage elements or using them as operands for single-precision arithmetic instructions. Values produced by single-precision load and arithmetic instructions are already single-precision values and can be stored directly into single format storage elements, or used directly as operands for single-precision arithmetic instructions, without preceding the store, or the arithmetic instruction, by a Floating Round to Single-Precision instruction.

## Programming Note

A single-precision value can be used in double-precision arithmetic operations. The reverse is true only if the double-precision value is representable in single format.
Some implementations may execute single-precision arithmetic instructions faster than double-precision arithmetic instructions. Therefore, if double-precision accuracy is not required, sin-gle-precision data and instructions should be used.

### 4.3.5.2 Integer-Valued Operands

Instructions are provided to round floating-point operands to integer values in floating-point format. To facilitate exchange of data between the floating-point and fixed-Point facilities, instructions are provided to convert between floating-point double format and fixed-point integer format in an FPR. Computation on integer-valued operands may be performed using arithmetic instructions of the required precision. (The results may not be integer values.) The two groups of instructions provided specifically to support integer-valued operands are described below.

## 1. Floating Round to Integer

The Floating Round to Integer instructions round a double-precision operand to an integer value in floating-point double format. These instructions may cause Invalid Operation (VXSNAN) exceptions. See Sections 4.3.6 and 4.5.1 for more information about rounding.

## 2. Floating Convert To/From Integer

The Floating Convert To Integer instructions convert a double-precision operand to a 32-bit or 64-bit signed fixed-point integer format. Variants are provided both to perform rounding based on
the value of FPSCR $R_{\text {RN }}$ and to round toward zero. These instructions may cause Invalid Operation (VXSNaN, VXCVI) and Inexact exceptions. The Floating Convert From Integer instruction converts a 64-bit signed fixed-point integer to a double-precision floating-point integer. Because of the limitations of the source format, only an Inexact exception may be generated.

### 4.3.6 Rounding

The material in this section applies to operations that have numeric operands (i.e., operands that are not infinities or NaNs ). Rounding the intermediate result of such an operation may cause an Overflow Exception, an Underflow Exception, or an Inexact Exception. The remainder of this section assumes that the operation causes no exceptions and that the result is numeric. See Section 4.3.2, "Value Representation" and Section 4.4, "Floating-Point Exceptions" for the cases not covered here.

The Arithmetic and Rounding and Conversion instructions round their intermediate results. With the exception of the Estimate instructions, these instructions produce an intermediate result that can be regarded as having infinite precision and unbounded exponent range. All but two groups of these instructions normalize or denormalize the intermediate result prior to rounding and then place the final result into the target FPR in double format. The Floating Round to Integer and Floating Convert To Integer instructions with biased exponents ranging from 1022 through 1074 are prepared for rounding by repetitively shifting the significand right one position and incrementing the biased exponent until it reaches a value of 1075. (Intermediate results with biased exponents 1075 or larger are already integers, and with biased exponents 1021 or less round to zero.) After rounding, the final result for Floating Round to Integer is normalized and put in double format, and for Floating Convert To Integer is converted to a signed fixed-point integer.

FPSCR bits FR and FI generally indicate the results of rounding. Each of the instructions which rounds its intermediate result sets these bits. If the fraction is incremented during rounding then FR is set to 1 , otherwise $F R$ is set to 0 . If the result is inexact then Fl is set to 1 , otherwise FI is set to zero. The Round to Integer instructions are exceptions to this rule, setting FR and FI to 0 . The Estimate instructions set FR and FI to undefined values. The remaining floating-point instructions do not alter FR and FI.

Four user-selectable rounding modes are provided through the Floating-Point Rounding Control field in the FPSCR. See Section 4.2.2, "Floating-Point Status and Control Register". These are encoded as follows.

## RN Rounding Mode

00 Round to Nearest
01 Round toward Zero
10 Round toward +Infinity
11 Round toward -Infinity
Let $Z$ be the intermediate arithmetic result or the operand of a convert operation. If $Z$ can be represented exactly in the target format, then the result in all rounding modes is $Z$ as represented in the target format. If $Z$ cannot be represented exactly in the target format, let $Z 1$ and $Z 2$ bound $Z$ as the next larger and next smaller numbers representable in the target format. Then Z1 or Z2 can be used to approximate the result in the target format.

Figure 55 shows the relation of $Z, Z 1$, and $Z 2$ in this case. The following rules specify the rounding in the four modes. "LSB" means "least significant bit".


Figure 55. Selection of Z1 and Z2

## Round to Nearest

Choose the value that is closer to $Z(Z 1$ or $Z 2)$. In case of a tie, choose the one that is even (least significant bit 0).

## Round toward Zero

 Choose the smaller in magnitude (Z1 or Z2).
## Round toward +Infinity

 Choose Z1.
## Round toward - Infinity

 Choose Z2.See Section 4.5.1, "Execution Model for IEEE Operations" on page 137 for a detailed explanation of rounding.

### 4.4 Floating-Point Exceptions

This architecture defines the following floating-point exceptions:

```
■ Invalid Operation Exception
    SNaN
    Infinity-Infinity
    Infinity\divInfinity
    Zero:Zero
    Infinity\timesZero
    Invalid Compare
    Software-Defined Condition
    Invalid Square Root
    Invalid Integer Convert
■ Zero Divide Exception
■ Overflow Exception
■ Underflow Exception
■ Inexact Exception
```

These exceptions, other than Invalid Operation Exception due to Software-Defined Condition, may occur during execution of computational instructions. An Invalid Operation Exception due to Software-Defined Condition occurs when a Move To FPSCR instruction sets FPSCR $_{\text {VXSOFT }}$ to 1.
Each floating-point exception, and each category of Invalid Operation Exception, has an exception bit in the FPSCR. In addition, each floating-point exception has a corresponding enable bit in the FPSCR. The exception bit indicates occurrence of the corresponding exception. If an exception occurs, the corresponding enable bit governs the result produced by the instruction and, in conjunction with the FEO and FE1 bits (see page 133), whether and how the system floating-point enabled exception error handler is invoked. (In general, the enabling specified by the enable bit is of invoking the system error handler, not of permitting the exception to occur. The occurrence of an exception depends only on the instruction and its inputs, not on the setting of any control bits. The only deviation from this general rule is that the occurrence of an Underflow Exception may depend on the setting of the enable bit.)
A single instruction, other than mtfsfi or mtfsf, may set more than one exception bit only in the following cases:

- Inexact Exception may be set with Overflow Exception.
- Inexact Exception may be set with Underflow Exception.
- Invalid Operation Exception ( SNaN ) is set with Invalid Operation Exception ( $\infty \times 0$ ) for Multiply-Add instructions for which the values being multiplied are infinity and zero and the value being added is an SNaN.
- Invalid Operation Exception (SNaN) may be set with Invalid Operation Exception (Invalid Compare) for Compare Ordered instructions.
- Invalid Operation Exception (SNaN) may be set with Invalid Operation Exception (Invalid Integer Convert) for Convert To Integer instructions.

When an exception occurs the writing of a result to the target register may be suppressed or a result may be delivered, depending on the exception.
The writing of a result to the target register is suppressed for the following kinds of exception, so that there is no possibility that one of the operands is lost:

## - Enabled Invalid Operation <br> - Enabled Zero Divide

For the remaining kinds of exception, a result is generated and written to the destination specified by the instruction causing the exception. The result may be a different value for the enabled and disabled conditions for some of these exceptions. The kinds of exception that deliver a result are the following:
■ Disabled Invalid Operation

- Disabled Zero Divide
- Disabled Overflow
- Disabled Underflow
- Disabled Inexact
- Enabled Overflow
- Enabled Underflow
- Enabled Inexact

Subsequent sections define each of the floating-point exceptions and specify the action that is taken when they are detected.

The IEEE standard specifies the handling of exceptional conditions in terms of "traps" and "trap handlers". In this architecture, an FPSCR exception enable bit of 1 causes generation of the result value specified in the IEEE standard for the "trap enabled" case; the expectation is that the exception will be detected by software, which will revise the result. An FPSCR exception enable bit of 0 causes generation of the "default result" value specified for the "trap disabled" (or "no trap occurs" or "trap is not implemented") case; the expectation is that the exception will not be detected by software, which will simply use the default result. The result to be delivered in each case for each exception is described in the sections below.

The IEEE default behavior when an exception occurs is to generate a default value and not to notify software. In this architecture, if the IEEE default behavior when an exception occurs is desired for all exceptions, all FPSCR exception enable bits should be set to 0 and Ignore Exceptions Mode (see below) should be used. In this case the system floating-point enabled exception error handler is not invoked, even if floating-point exceptions occur: software can inspect the FPSCR exception bits if necessary, to determine whether exceptions have occurred.

In this architecture, if software is to be notified that a given kind of exception has occurred, the corresponding FPSCR exception enable bit must be set to 1 and a mode other than Ignore Exceptions Mode must be used. In this case the system floating-point enabled exception error handler is invoked if an enabled float-
ing-point exception occurs. The system floating-point enabled exception error handler is also invoked if a Move To FPSCR instruction causes an exception bit and the corresponding enable bit both to be 1 ; the Move To FPSCR instruction is considered to cause the enabled exception.

The FE0 and FE1 bits control whether and how the system floating-point enabled exception error handler is invoked if an enabled floating-point exception occurs. The location of these bits and the requirements for altering them are described in Book III. (The system floating-point enabled exception error handler is never invoked because of a disabled floating-point exception.) The effects of the four possible settings of these bits are as follows.

## FE0 FE1 Description

00 Ignore Exceptions Mode
Floating-point exceptions do not cause the system floating-point enabled exception error handler to be invoked.
01 Imprecise Nonrecoverable Mode
The system floating-point enabled exception error handler is invoked at some point at or beyond the instruction that caused the enabled exception. It may not be possible to identify the excepting instruction or the data that caused the exception. Results produced by the excepting instruction may have been used by or may have affected subsequent instructions that are executed before the error handler is invoked.
10 Imprecise Recoverable Mode
The system floating-point enabled exception error handler is invoked at some point at or beyond the instruction that caused the enabled exception. Sufficient information is provided to the error handler that it can identify the excepting instruction and the operands, and correct the result. No results produced by the excepting instruction have been used by or have affected subsequent instructions that are executed before the error handler is invoked.
11 Precise Mode
The system floating-point enabled exception error handler is invoked precisely at the instruction that caused the enabled exception.

In all cases, the question of whether a floating-point result is stored, and what value is stored, is governed by the FPSCR exception enable bits, as described in subsequent sections, and is not affected by the value of the FE0 and FE1 bits.

In all cases in which the system floating-point enabled exception error handler is invoked, all instructions
before the instruction at which the system floating-point enabled exception error handler is invoked have completed, and no instruction after the instruction at which the system floating-point enabled exception error handler is invoked has begun execution. The instruction at which the system floating-point enabled exception error handler is invoked has completed if it is the excepting instruction and there is only one such instruction. Otherwise it has not begun execution (or may have been partially executed in some cases, as described in Book III).

## Programming Note

In any of the three non-Precise modes, a Float-ing-Point Status and Control Register instruction can be used to force any exceptions, due to instructions initiated before the Floating-Point Status and Control Register instruction, to be recorded in the FPSCR. (This forcing is superfluous for Precise Mode.)
In either of the Imprecise modes, a Floating-Point Status and Control Register instruction can be used to force any invocations of the system floating-point enabled exception error handler, due to instructions initiated before the Floating-Point Status and Control Register instruction, to occur. (This forcing has no effect in Ignore Exceptions Mode, and is superfluous for Precise Mode.)

The last sentence of the paragraph preceding this Programming Note can apply only in the Imprecise modes, or if the mode has just been changed from Ignore Exceptions Mode to some other mode. (It always applies in the latter case.)

In order to obtain the best performance across the widest range of implementations, the programmer should obey the following guidelines.

- If the IEEE default results are acceptable to the application, Ignore Exceptions Mode should be used with all FPSCR exception enable bits set to 0 .
- If the IEEE default results are not acceptable to the application, Imprecise Nonrecoverable Mode should be used, or Imprecise Recoverable Mode if recoverability is needed, with FPSCR exception enable bits set to 1 for those exceptions for which the system floating-point enabled exception error handler is to be invoked.
■ Ignore Exceptions Mode should not, in general, be used when any FPSCR exception enable bits are set to 1 .
- Precise Mode may degrade performance in some implementations, perhaps substantially, and therefore should be used only for debugging and other specialized applications.


### 4.4.1 Invalid Operation Exception

### 4.4.1.1 Definition

An Invalid Operation Exception occurs when an operand is invalid for the specified operation. The invalid operations are:

- Any floating-point operation on a Signaling NaN (SNaN)
- For add or subtract operations, magnitude subtraction of infinities $(\infty-\infty)$
■ Division of infinity by infinity ( $\infty \div \infty$ )
- Division of zero by zero ( $0 \div 0$ )
- Multiplication of infinity by zero ( $\infty \times 0$ )
- Ordered comparison involving a NaN (Invalid Compare)
- Square root or reciprocal square root of a negative (and nonzero) number (Invalid Square Root)
- Integer convert involving a number too large in magnitude to be represented in the target format, or involving an infinity or a NaN (Invalid Integer Convert)

An Invalid Operation Exception also occurs when an $\boldsymbol{m t f s f i}, \boldsymbol{m t f s f}$, or mtfsb1 instruction is executed that sets FPSCR $_{\text {VXSOFT }}$ to 1 (Software-Defined Condition).

### 4.4.1.2 Action

The action to be taken depends on the setting of the Invalid Operation Exception Enable bit of the FPSCR.

When Invalid Operation Exception is enabled ( $\mathrm{FPSCR}_{\mathrm{VE}}=1$ ) and an Invalid Operation Exception occurs, the following actions are taken:

1. One or two Invalid Operation Exceptions are set

| FPSCR $_{V X S N A N}$ | (if SNaN) |
| :--- | ---: |
| FPSCR $_{V X I S I}$ | (if $\infty-\infty$ ) |
| FPSCR $_{V X I D I}$ | (if $\infty \div \infty$ ) |
| FPSCR $_{V X Z D Z}$ | (if $0 \div 0$ ) |
| FPSCR $_{V X I M Z ~}$ | (if $\infty \times 0$ ) |

FPSCR
(if invalid comp)
FPSCR ${ }_{\text {VXSOFT }}$
(if sfw-def cond)
(if invalid sqrt)
(if invalid int cvrt)
2. If the operation is an arithmetic, Floating Round to Single-Precision, Floating Round to Integer, or convert to integer operation,
the target FPR is unchanged
FPSCR $_{\text {FR FI }}$ are set to zero
FPSCR $_{\text {FPRF }}$ is unchanged
3. If the operation is a compare,
$\mathrm{FPSCR}_{\text {FR FI C }}$ are unchanged
FPSCR $_{\text {FPCC }}$ is set to reflect unordered
4. If an mtfsfi, mtfsf, or mtfsb1 instruction is executed that sets FPSCR ${ }_{\text {VXSOFT }}$ to 1,

The FPSCR is set as specified in the instruction description.

When Invalid Operation Exception is disabled ( $\mathrm{FPSCR}_{\mathrm{VE}}=0$ ) and an Invalid Operation Exception occurs, the following actions are taken:

1. One or two Invalid Operation Exceptions are set

| FPSCR $_{V X S N A N}$ | (if SNaN) |
| :--- | ---: |
| FPSCR $_{V X I S I}$ | (if $\infty-\infty$ ) |
| FPSCR $_{V X I D I}$ | (if $\infty \div \infty$ ) |
| FPSCR $_{V X Z D Z}$ | (if $0 \div 0$ ) |
| FPSCR | (if $\infty \times 0$ ) |

FPSCR $_{V \times V \mathrm{~V}}$ (if invalid comp)
FPSCR $_{\text {VXSOFT }}$ (if sfw-def cond)
FPSCR ${ }_{V x \text { xart }}$
(if invalid sqrt)
FPSCR $_{V x C V I}$
(if invalid int cvrt)
2. If the operation is an arithmetic or Floating Round to Single-Precision operation,
the target FPR is set to a Quiet NaN
FPSCR $_{\text {FR FI }}$ are set to zero
FPSCR $_{\text {FPRF }}$ is set to indicate the class of the result (Quiet NaN)
3. If the operation is a convert to 64-bit integer operation,
the target FPR is set as follows:
FRT is set to the most positive 64-bit integer if the operand in FRB is a positive number or $+\infty$, and to the most negative 64-bit integer if the operand in FRB is a negative number, $-\infty$, or NaN
FPSCR $_{\text {FR FI }}$ are set to zero
FPSCR $_{\text {FPRF }}$ is undefined
4. If the operation is a convert to 32-bit integer operation,
the target FPR is set as follows:
$\mathrm{FRT}_{0: 31} \leftarrow$ undefined
$\mathrm{FRT}_{32: 63}$ are set to the most positive 32-bit integer if the operand in FRB is a positive number or +infinity, and to the most negative 32-bit integer if the operand in FRB is a negative number, -infinity, or NaN
FPSCR $_{\text {FR FI }}$ are set to zero
FPSCR $_{\text {FPRF }}$ is undefined
5. If the operation is a compare,

FPSCR $_{\text {FR FI }}$ are unchanged
FPSCR $_{\text {FPCC }}$ is set to reflect unordered
6. If an mtfsfi, mtfsf, or mtfsb1 instruction is executed that sets FPSCR $_{\text {VXSOFT }}$ to 1,

The FPSCR is set as specified in the instruction description.

### 4.4.2 Zero Divide Exception

### 4.4.2.1 Definition

A Zero Divide Exception occurs when a Divide instruction is executed with a zero divisor value and a finite nonzero dividend value. It also occurs when a Reciprocal Estimate instruction (fre[s] or frsqrte[s]) is executed with an operand value of zero.

### 4.4.2.2 Action

The action to be taken depends on the setting of the Zero Divide Exception Enable bit of the FPSCR.
When Zero Divide Exception is enabled (FPSCR ${ }_{\text {ZE }}=1$ ) and a Zero Divide Exception occurs, the following actions are taken:

1. Zero Divide Exception is set
$\mathrm{FPSCR}_{Z X} \leftarrow 1$
2. The target FPR is unchanged
3. $\mathrm{FPSCR}_{\text {FR FI }}$ are set to zero
4. FPSCR FPRF is unchanged

When Zero Divide Exception is disabled ( $\mathrm{FPSCR}_{\mathrm{ZE}}=0$ ) and a Zero Divide Exception occurs, the following actions are taken:

1. Zero Divide Exception is set $\mathrm{FPSCR}_{Z x} \leftarrow 1$
2. The target FPR is set to $\pm$ Infinity, where the sign is determined by the XOR of the signs of the operands
3. FPSCR $_{\text {FR FI }}$ are set to zero
4. FPSCR $_{\text {FPRF }}$ is set to indicate the class and sign of the result ( $\pm$ Infinity)

### 4.4.3 Overflow Exception

### 4.4.3.1 Definition

An Overflow Exception occurs when the magnitude of what would have been the rounded result if the exponent range were unbounded exceeds that of the largest finite number of the specified result precision.

### 4.4.3.2 Action

The action to be taken depends on the setting of the Overflow Exception Enable bit of the FPSCR.

When Overflow Exception is enabled (FPSCR ${ }_{\mathrm{OE}}=1$ ) and an Overflow Exception occurs, the following actions are taken:

1. Overflow Exception is set

FPSCR $_{\mathrm{OX}} \leftarrow 1$
2. For double-precision arithmetic instructions, the exponent of the normalized intermediate result is adjusted by subtracting 1536
3. For single-precision arithmetic instructions and the Floating Round to Single-Precision instruction, the exponent of the normalized intermediate result is adjusted by subtracting 192
4. The adjusted rounded result is placed into the target FPR
5. FPSCR $_{\text {FPRF }}$ is set to indicate the class and sign of the result ( $\pm$ Normal Number)
When Overflow Exception is disabled (FPSCR $\mathrm{OE}^{=}=0$ ) and an Overflow Exception occurs, the following actions are taken:

1. Overflow Exception is set
$\mathrm{FPSCR}_{\mathrm{OX}} \leftarrow 1$
2. Inexact Exception is set

FPSCR $_{X X} \leftarrow 1$
3. The result is determined by the rounding mode ( $\mathrm{FPSCR}_{\mathrm{RN}}$ ) and the sign of the intermediate result as follows:

- Round to Nearest

Store $\pm$ Infinity, where the sign is the sign of the intermediate result

- Round toward Zero Store the format's largest finite number with the sign of the intermediate result
- Round toward + Infinity

For negative overflow, store the format's most negative finite number; for positive overflow, store +Infinity

- Round toward - Infinity

For negative overflow, store - Infinity; for positive overflow, store the format's largest finite number
4. The result is placed into the target FPR
5. FPSCR $_{F R}$ is undefined
6. FPSCR $_{\text {FI }}$ is set to 1
7. FPSCR $_{\text {FPRF }}$ is set to indicate the class and sign of the result ( $\pm$ Infinity or $\pm$ Normal Number)

### 4.4.4 Underflow Exception

### 4.4.4.1 Definition

Underflow Exception is defined separately for the enabled and disabled states:

- Enabled:

Underflow occurs when the intermediate result is "Tiny".

- Disabled:

Underflow occurs when the intermediate result is "Tiny" and there is "Loss of Accuracy".
A "Tiny" result is detected before rounding, when a nonzero intermediate result computed as though both the precision and the exponent range were unbounded would be less in magnitude than the smallest normalized number.

If the intermediate result is "Tiny" and Underflow Exception is disabled (FPSCR ${ }_{U E}=0$ ) then the intermediate result is denormalized (see Section 4.3.4, "Normalization and Denormalization" on page 129) and rounded (see Section 4.3.6, "Rounding" on page 131) before being placed into the target FPR.
"Loss of Accuracy" is detected when the delivered result value differs from what would have been computed were both the precision and the exponent range unbounded.

### 4.4.4.2 Action

The action to be taken depends on the setting of the Underflow Exception Enable bit of the FPSCR.

When Underflow Exception is enabled (FPSCR ${ }_{\text {UE }}=1$ ) and an Underflow Exception occurs, the following actions are taken:

1. Underflow Exception is set

FPSCR $_{U X} \leftarrow 1$
2. For double-precision arithmetic instructions, the exponent of the normalized intermediate result is adjusted by adding 1536
3. For single-precision arithmetic instructions and the Floating Round to Single-Precision instruction, the exponent of the normalized intermediate result is adjusted by adding 192
4. The adjusted rounded result is placed into the target FPR
5. FPSCR FPRF is set to indicate the class and sign of the result ( $\pm$ Normalized Number)

## Programming Note

The FR and FI bits are provided to allow the system floating-point enabled exception error handler, when invoked because of an Underflow Exception, to simulate a "trap disabled" environment. That is, the FR and FI bits allow the system floating-point enabled exception error handler to unround the result, thus allowing the result to be denormalized.

When Underflow Exception is disabled (FPSCR ${ }_{U E}=0$ ) and an Underflow Exception occurs, the following actions are taken:

1. Underflow Exception is set

$$
\mathrm{FPSCR}_{U X} \leftarrow 1
$$

2. The rounded result is placed into the target FPR
3. FPSCR $_{\text {FPRF }}$ is set to indicate the class and sign of the result ( $\pm$ Normalized Number, $\pm$ Denormalized Number, or $\pm$ Zero)

### 4.4.5 Inexact Exception

### 4.4.5.1 Definition

An Inexact Exception occurs when one of two conditions occur during rounding:

1. The rounded result differs from the intermediate result assuming both the precision and the exponent range of the intermediate result to be unbounded. In this case the result is said to be inexact. (If the rounding causes an enabled Overflow Exception or an enabled Underflow Exception, an Inexact Exception also occurs only if the significands of the rounded result and the intermediate result differ.)
2. The rounded result overflows and Overflow Exception is disabled.

### 4.4.5.2 Action

The action to be taken does not depend on the setting of the Inexact Exception Enable bit of the FPSCR.

When an Inexact Exception occurs, the following actions are taken:

1. Inexact Exception is set

FPSCR $_{X X} \leftarrow 1$
2. The rounded or overflowed result is placed into the target FPR
3. FPSCR $_{\text {FPRF }}$ is set to indicate the class and sign of the result

## Programming Note

In some implementations, enabling Inexact Exceptions may degrade performance more than does enabling other types of floating-point exception.

### 4.5 Floating-Point Execution Models

All implementations of this architecture must provide the equivalent of the following execution models to ensure that identical results are obtained.

Special rules are provided in the definition of the computational instructions for the infinities, denormalized numbers and NaNs. The material in the remainder of this section applies to instructions that have numeric operands and a numeric result (i.e., operands and result that are not infinities or NaNs ), and that cause no exceptions. See Section 4.3.2 and Section 4.4 for the cases not covered here.

Although the double format specifies an 11-bit exponent, exponent arithmetic makes use of two additional bits to avoid potential transient overflow conditions. One extra bit is required when denormalized dou-ble-precision numbers are prenormalized. The second bit is required to permit the computation of the adjusted exponent value in the following cases when the corresponding exception enable bit is 1 :
■ Underflow during multiplication using a denormalized operand.
■ Overflow during division using a denormalized divisor.
The IEEE standard includes 32-bit and 64-bit arithmetic. The standard requires that single-precision arithmetic be provided for single-precision operands. The standard permits double-precision floating-point operations to have either (or both) single-precision or dou-ble-precision operands, but states that single-precision floating-point operations should not accept double-precision operands. The Power ISA follows these guidelines; double-precision arithmetic instructions can have operands of either or both precisions, while single-precision arithmetic instructions require all operands to be single-precision. Double-precision arithmetic instructions and fcfid produce double-precision values, while single-precision arithmetic instructions produce sin-gle-precision values.
For arithmetic instructions, conversions from dou-ble-precision to single-precision must be done explicitly by software, while conversions from single-precision to double-precision are done implicitly.

### 4.5.1 Execution Model for IEEE Operations

The following description uses 64-bit arithmetic as an example. 32-bit arithmetic is similar except that the FRACTION is a 23 -bit field, and the single-precision Guard, Round, and Sticky bits (described in this section) are logically adjacent to the 23-bit FRACTION field.

IEEE-conforming significand arithmetic is considered to be performed with a floating-point accumulator having the following format, where bits 0:55 comprise the significand of the intermediate result.


Figure 56. IEEE 64-bit execution model
The $S$ bit is the sign bit.
The C bit is the carry bit, which captures the carry out of the significand.
The $L$ bit is the leading unit bit of the significand, which receives the implicit bit from the operand.

The FRACTION is a 52 -bit field that accepts the fraction of the operand.
The Guard (G), Round (R), and Sticky (X) bits are extensions to the low-order bits of the accumulator. The $G$ and $R$ bits are required for postnormalization of the result. The G, R, and $X$ bits are required during rounding to determine if the intermediate result is equally near the two nearest representable values. The $X$ bit serves as an extension to the $G$ and $R$ bits by representing the logical OR of all bits that may appear to the low-order side of the R bit, due either to shifting the accumulator right or to other generation of low-order result bits. The $G$ and $R$ bits participate in the left shifts with zeros being shifted into the R bit. Figure 57 shows the significance of the $G, R$, and $X$ bits with respect to the intermediate result (IR), the representable number next lower in magnitude (NL), and the representable number next higher in magnitude (NH).

| G R X | Interpretation |
| :---: | :---: |
| 000 | IR is exact |
| 001 |  |
| 010 | IR closer to NL |
| 011 |  |
| 100 | IR midway between NL and NH |
| 101 |  |
| 110 | IR closer to NH |
| 111 |  |

Figure 57. Interpretation of $G, R$, and $X$ bits
Figure 58 shows the positions of the Guard, Round, and Sticky bits for double-precision and single-precision floating-point numbers relative to the accumulator illustrated in Figure 56.

| Format | Guard | Round | Sticky |
| :--- | :--- | :--- | :--- |
| Double | G bit | R bit | X bit |
| Single | 24 | 25 | OR of 26:52, G, R, X |

Figure 58. Location of the Guard, Round, and Sticky bits in the IEEE execution model

## Version 3.0

The significand of the intermediate result is prepared for rounding by shifting its contents right, if required, until the least significant bit to be retained is in the low-order bit position of the fraction. Four user-selectable rounding modes are provided through FPSCR ${ }_{\text {RN }}$ as described in Section 4.3.6, "Rounding" on page 131. Using Z1 and Z2 as defined on page 131, the rules for rounding in each mode are as follows.

## - Round to Nearest

## Guard bit = 0

The result is truncated. (Result exact (GRX=000) or closest to next lower value in magnitude (GRX=001, 010, or 011))

## Guard bit = 1

Depends on Round and Sticky bits:

## Case a

If the Round or Sticky bit is 1 (inclusive), the result is incremented. (Result closest to next higher value in magnitude ( $G R X=101,110$, or 111))

## Case b

If the Round and Sticky bits are 0 (result midway between closest representable values), then if the low-order bit of the result is 1 the result is incremented. Otherwise (the low-order bit of the result is 0 ) the result is truncated (this is the case of a tie rounded to even).

## - Round toward Zero

Choose the smaller in magnitude of Z1 or Z2. If the Guard, Round, or Sticky bit is nonzero, the result is inexact.

■ Round toward + Infinity
Choose Z1.
■ Round toward - Infinity
Choose Z2.
If rounding results in a carry into C , the significand is shifted right one position and the exponent is incremented by one. This yields an inexact result, and possibly also exponent overflow. If any of the Guard, Round, or Sticky bits is nonzero, then the result is also inexact. Fraction bits are stored to the target FPR. For Floating Round to Integer, Floating Round to Single-Precision, and single-precision arithmetic instructions, low-order zeros must be appended as appropriate to fill out the double-precision fraction.

### 4.5.2 Execution Model for Multiply-Add Type Instructions

The Power ISA provides a special form of instruction that performs up to three operations in one instruction (a multiplication, an addition, and a negation). With this added capability comes the special ability to produce a more exact intermediate result as input to the rounder. 32-bit arithmetic is similar except that the FRACTION field is smaller.

Multiply-add significand arithmetic is considered to be performed with a floating-point accumulator having the following format, where bits 0:106 comprise the significand of the intermediate result.


Figure 59. Multiply-add 64-bit execution model
The first part of the operation is a multiplication. The multiplication has two 53 -bit significands as inputs, which are assumed to be prenormalized, and produces a result conforming to the above model. If there is a carry out of the significand (into the C bit), then the significand is shifted right one position, shifting the $L$ bit (leading unit bit) into the most significant bit of the FRACTION and shifting the C bit (carry out) into the L bit. All 106 bits (L bit, the FRACTION) of the product take part in the add operation. If the exponents of the two inputs to the adder are not equal, the significand of the operand with the smaller exponent is aligned (shifted) to the right by an amount that is added to that exponent to make it equal to the other input's exponent. Zeros are shifted into the left of the significand as it is aligned and bits shifted out of bit 105 of the significand are ORed into the X ' bit. The add operation also produces a result conforming to the above model with the $X$ ' bit taking part in the add operation.

The result of the addition is then normalized, with all bits of the addition result, except the $X^{\prime}$ bit, participating in the shift. The normalized result serves as the intermediate result that is input to the rounder.

For rounding, the conceptual Guard, Round, and Sticky bits are defined in terms of accumulator bits. Figure 60 shows the positions of the Guard, Round, and Sticky bits for double-precision and single-precision float-ing-point numbers in the multiply-add execution model.

| Format | Guard | Round | Sticky |
| :--- | :--- | :--- | :--- |
| Double | 53 | 54 | OR of $55: 105, X^{\prime}$ |
| Single | 24 | 25 | OR of $26: 105, X^{\prime}$ |

The rules for rounding the intermediate result are the same as those given in Section 4.5.1.

## Figure 60. Location of the Guard, Round, and <br> Sticky bits in the multiply-add execution model

If the instruction is Floating Negative Multiply-Add or Floating Negative Multiply-Subtract, the final result is negated.

### 4.6 Floating-Point Facility Instructions

### 4.6.1 Floating-Point Storage Access Instructions

The Storage Access instructions compute the effective address (EA) of the storage to be accessed as described in Section 1.11.3, "Effective Address Calculation" on page 27.

## Programming Note

The la extended mnemonic permits computing an effective address as a Load or Store instruction would, but loads the address itself into a GPR rather than loading the value that is in storage at that address. This extended mnemonic is described in Section C.10, "Miscellaneous Mnemonics" on page 806.

### 4.6.1.1 Storage Access Exceptions

Storage accesses will cause the system data storage error handler to be invoked if the program is not allowed to modify the target storage (Store only), or if the program attempts to access storage that is unavailable.

### 4.6.2 Floating-Point Load Instructions

There are three basic forms of load instruction: sin-gle-precision, double-precision, and integer. The integer form is provided by the Load Floating-Point as Integer Word Algebraic instruction, described on page 144. Because the FPRs support only float-ing-point double format, single-precision Load Float-ing-Point instructions convert single-precision data to double format prior to loading the operand into the target FPR. The conversion and loading steps are as follows.

Let WORD $_{0: 31}$ be the floating-point single-precision operand accessed from storage.

```
Normalized Operand
if WORD \(_{1: 8}>0\) and WORD \(_{1: 8}<255\) then
    \(\mathrm{FRT}_{0: 1} \leftarrow\) WORD \(_{0: 1}\)
    \(\mathrm{FRT}_{2} \leftarrow \neg \mathrm{WORD}_{1}\)
    \(\mathrm{FRT}_{3} \leftarrow \neg \mathrm{WORD}_{1}\)
    \(\mathrm{FRT}_{4} \leftarrow \neg \mathrm{WORD}_{1}\)
    FRT \(_{5: 63} \leftarrow\) WORD \(_{2: 31} \|^{29} 0\)
Denormalized Operand
    if WORD \(_{1: 8}=0\) and WORD \(_{9: 31} \neq 0\) then
    sign \(\leftarrow\) WORD \(_{0}\)
    \(\exp \leftarrow-126\)
    frac \(_{0: 52} \leftarrow 0 \mathrm{bO} \|\) WORD \(_{9: 31} \|{ }^{29} 0\)
    normalize the operand
        do while frac \({ }_{0}=0\)
            \(\mathrm{frac}_{0: 52} \leftarrow \mathrm{frac}_{1: 52}\) II Ob0
```

$$
\begin{aligned}
& \quad \exp \leftarrow \exp -1 \\
& \text { FRT }_{0} \leftarrow \operatorname{sign} \\
& \text { FRT }_{1: 11} \leftarrow \exp +1023 \\
& \text { FRT }_{12: 63} \leftarrow \text { frac }_{1: 52}
\end{aligned}
$$

## Zero / Infinity / NaN

if WORD $_{1: 8}=255$ or WORD $_{1: 31}=0$ then
$\mathrm{FRT}_{0: 1} \leftarrow$ WORD $_{0: 1}$
$\mathrm{FRT}_{2} \leftarrow \mathrm{WORD}_{1}$
$\mathrm{FRT}_{3} \leftarrow \mathrm{WORD}_{1}$
$\mathrm{FRT}_{4} \leftarrow \mathrm{WORD}_{1}$
$\mathrm{FRT}_{5: 63} \leftarrow$ WORD $_{2: 31}{ }^{1 / 29} 0$
For double-precision Load Floating-Point instructions and for the Load Floating-Point as Integer Word Algebraic instruction no conversion is required, as the data from storage are copied directly into the FPR.

Many of the Load Floating-Point instructions have an "update" form, in which register RA is updated with the effective address. For these forms, if $R A \neq 0$, the effective address is placed into register RA and the storage element (word or doubleword) addressed by EA is loaded into FRT.

Note: Recall that RA and RB denote General Purpose Registers, while FRT denotes a Floating-Point Register.

## Load Floating-Point Single D-form

Ifs $\quad$ FRT, $D(R A)$

| 48 | FRT | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  | 3 |

```
if RA = 0 then b }\leftarrow
else b}\leftarrow(RA
EA \leftarrow b + EXTS(D)
FRT \leftarrow DOUBLE(MEM(EA, 4))
```

Let the effective address (EA) be the sum (RA|O)+D.
The word in storage addressed by EA is interpreted as a floating-point single-precision operand. This word is converted to floating-point double format (see page 141) and placed into register FRT.

## Special Registers Altered: <br> None <br> Load Floating-Point Single with Update D-form

Ifsu

| 49 | FRT,D(RA) |  |  |  |  |
| ---: | ---: | ---: | ---: | ---: | ---: |
| 0 | 6 | RA |  |  | D |

```
EA \leftarrow(RA) + EXTS(D)
FRT \leftarrow DOUBLE(MEM(EA, 4))
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+D.
The word in storage addressed by EA is interpreted as a floating-point single-precision operand. This word is converted to floating-point double format (see page 141) and placed into register FRT.

EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None

## Load Floating-Point Single Indexed X-form

| Ifsx FRT,RA,RB |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{\|ll}  & 31 \\ 0 & \end{array}$ | ${ }_{6}$ FRT | $11{ }^{R}$ | ${ }_{16} \mathrm{RB}$ | $21535$ | 1 31 |

```
if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(\mathrm{ RA)
EA \leftarrow b + (RB)
FRT \leftarrow DOUBLE (MEM (EA, 4))
```

Let the effective address (EA) be the sum (RA|0)+(RB).
The word in storage addressed by EA is interpreted as a floating-point single-precision operand. This word is converted to floating-point double format (see page 141) and placed into register FRT.

## Special Registers Altered:

None

## Load Floating-Point Single with Update Indexed X-form

Ifsux FRT,RA,RB

| 31 | FRT | RA | RB |  | 567 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |

```
EA \leftarrow(RA) + (RB)
FRT}\leftarrow\operatorname{DOUBLE(MEM(EA, 4))
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+(RB).
The word in storage addressed by EA is interpreted as a floating-point single-precision operand. This word is converted to floating-point double format (see page 141) and placed into register FRT.

EA is placed into register RA.
If $R A=0$, the instruction form is invalid.

## Special Registers Altered:

None

## Load Floating-Point Double D-form

## Ifd $\quad \mathrm{FRT}, \mathrm{D}(\mathrm{RA})$

| 50 | FRT | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  | 3 |

```
if RA = 0 then b }\leftarrow
else b
EA\leftarrowb+EXTS(D)
FRT \leftarrow MEM(EA, 8)
```

Let the effective address (EA) be the sum (RA|0)+D.
The doubleword in storage addressed by EA is loaded into register FRT.
Special Registers Altered:
None

## Load Floating-Point Double with Update D-form

Ifdu \begin{tabular}{l}
FRT, D(RA) <br>

| 51 | FRT | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  | 31 |

\end{tabular}.

```
EA \leftarrow(RA) + EXTS (D)
FRT \leftarrow MEM(EA, 8)
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+D.
The doubleword in storage addressed by EA is loaded into register FRT.

EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None

## Load Floating-Point Double Indexed X-form

Ifdx
FRT,RA,RB

| 31 | FRT | RA | RB |  | 599 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

```
if RA = 0 then b }\leftarrow
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
FRT \leftarrow MEM(EA, 8)
```

Let the effective address (EA) be the sum (RA|0)+(RB).
The doubleword in storage addressed by EA is loaded into register FRT.

Special Registers Altered:
None

## Load Floating-Point Double with Update Indexed X-form

Ifdux FRT,RA,RB

| 31 | FRT | RA | RB |  | 631 | 1 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 |  | 16 |  |

```
EA}\leftarrow(RA)+(RB
FRT \leftarrowMEM(EA, 8)
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+(RB).
The doubleword in storage addressed by EA is loaded into register FRT.

EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None

## Version 3.0

## Load Floating-Point as Integer Word <br> Algebraic Indexed X-form

Ifiwax FRT,RA,RB

| 31 | FRT | RA | RB |  | 855 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 16 |  |  |  |  |  |

```
if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(RA
EA \leftarrow b + (RB)
FRT \leftarrow EXTS(MEM(EA, 4))
```

Let the effective address (EA) be the sum (RA|0)+(RB).
The word in storage addressed by EA is loaded into $\mathrm{FRT}_{32: 63} . \mathrm{FRT}_{0: 31}$ are filled with a copy of bit 0 of the loaded word.

## Special Registers Altered:

None

## Load Floating-Point as Integer Word and

 Zero Indexed X-formIfiwzx FRT,RA,RB

| 31 | FRT | RA | RB |  | 887 | 1 <br> 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |

```
if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(\textrm{RA}
EA \leftarrow b + (RB)
FRT \leftarrow '320 || MEM(EA, 4)
```

Let the effective address (EA) be the sum (RA|0)+(RB).
The word in storage addressed by EA is loaded into $\mathrm{FRT}_{32: 63} . \mathrm{FRT}_{0: 31}$ are set to 0 .

## Special Registers Altered:

None

### 4.6.3 Floating-Point Store Instructions

There are three basic forms of store instruction: sin-gle-precision, double-precision, and integer. The integer form is provided by the Store Floating-Point as Integer Word instruction, described on page 148. Because the FPRs support only floating-point double format for floating-point data, single-precision Store Floating-Point instructions convert double-precision data to single format prior to storing the operand into storage. The conversion steps are as follows.

Let $\mathrm{WORD}_{0: 31}$ be the word in storage written to.
No Denormalization Required (includes Zero / Infinity / NaN)
if $\mathrm{FRS}_{1: 11}>896$ or $\mathrm{FRS}_{1: 63}=0$ then
WORD $_{0: 1} \leftarrow \mathrm{FRS}_{0: 1}$
WORD $_{2: 31} \leftarrow \mathrm{FRS}_{5: 34}$
Denormalization Required
if $874 \leq \mathrm{FRS}_{1: 11} \leq 896$ then
sign $\leftarrow \mathrm{FRS}_{0}$
$\exp \leftarrow \mathrm{FRS}_{1: 11}-1023$
frac $_{0: 52} \leftarrow 0 \mathrm{~b} 1$ II $\mathrm{FRS}_{12: 63}$
denormalize operand
do while exp <-126
frac $_{0: 52} \leftarrow 0 \mathrm{ODO}$ II frac $\mathrm{f}_{0: 51}$
$\exp \leftarrow \exp +1$
WORD $_{0} \leftarrow$ sign
WORD $_{1: 8} \leftarrow 0 \times 00$
WORD $_{9: 31} \leftarrow$ frac $_{1: 23}$
else WORD $\leftarrow$ undefined
Notice that if the value to be stored by a single-precision Store Floating-Point instruction is larger in magnitude than the maximum number representable in single format, the first case above (No Denormalization Required) applies. The result stored in WORD is then a well-defined value, but is not numerically equal to the value in the source register (i.e., the result of a sin-
gle-precision Load Floating-Point from WORD will not compare equal to the contents of the original source register).

For double-precision Store Floating-Point instructions and for the Store Floating-Point as Integer Word instruction no conversion is required, as the data from the FPR are copied directly into storage.

Many of the Store Floating-Point instructions have an "update" form, in which register RA is updated with the effective address. For these forms, if $R A \neq 0$, the effective address is placed into register RA.

Note: Recall that RA and RB denote General Purpose Registers, while FRS denotes a Floating-Point Register.

# Store Floating-Point Single D-form 

stfs $\quad$ FRS,D(RA)

| 52 | FRS | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  | 31 |

```
if RA = 0 then b \leftarrow 
else b}\leftarrow(RA
EA}\leftarrow\textrm{b}+\operatorname{EXTS}(\textrm{D}
MEM(EA, 4) \leftarrow SINGLE((FRS))
```

Let the effective address (EA) be the sum (RA|0)+D.
The contents of register FRS are converted to single format (see page 145) and stored into the word in storage addressed by EA.

## Special Registers Altered:

None

## Store Floating-Point Single with Update D-form

stfsu $\quad$ FRS, D(RA)

| 53 | FRS | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |

```
EA \leftarrow(RA) + EXTS (D)
MEM(EA, 4) \leftarrow SINGLE((FRS))
RA}\leftarrowE
```

Let the effective address (EA) be the sum (RA)+D.
The contents of register FRS are converted to single format (see page 145) and stored into the word in storage addressed by EA.
EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None

## Store Floating-Point Single Indexed X-form

```
stfsx FRS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & FRS & RA & RB & \multicolumn{2}{|c|}{663} & \(/\) \\
0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

```
if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(\textrm{RA
EA}\leftarrow\textrm{b}+(\textrm{RB}
MEM(EA, 4)}\leftarrow\operatorname{SINGLE((FRS))
```

Let the effective address (EA) be the sum (RA|0)+(RB).
The contents of register FRS are converted to single format (see page 145) and stored into the word in storage addressed by EA.

## Special Registers Altered:

None

## Store Floating-Point Single with Update Indexed X-form

```
stfsux FRS,RA,RB
```

| 31 | FRS | RA | RB |  | 695 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |

```
EA}\leftarrow(RA)+(RB
MEM(EA, 4) \leftarrow SINGLE((FRS))
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+(RB).
The contents of register FRS are converted to single format (see page 145) and stored into the word in storage addressed by EA.
EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None

## Store Floating-Point Double D-form

stfd $\quad$ FRS,D(RA)

| 54 | FRS | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  | 3 |

```
if RA = 0 then b \leftarrow 
else b b (RA)
EA \leftarrow b + EXTS(D)
MEM(EA, 8) \leftarrow(FRS)
```

Let the effective address (EA) be the sum (RA|0)+D.
The contents of register FRS are stored into the doubleword in storage addressed by EA.

Special Registers Altered:
None

## Store Floating-Point Double with Update

 D-formstfdu
FRS,D(RA)

| 55 | FRS | RA |  | D | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 |  | 3 |

```
EA \leftarrow(RA) + EXTS(D)
MEM(EA, 8) \leftarrow(FRS)
RA}\leftarrowE
```

Let the effective address (EA) be the sum (RA)+D.
The contents of register FRS are stored into the doubleword in storage addressed by EA.

EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None

## Store Floating-Point Double Indexed X-form

| stfdx FRS,RA,RB |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $31$ | ${ }_{6} \text { FRS }$ | $\int_{11} \mathrm{RA}$ | ${ }_{16} \mathrm{RB}$ | $21$ | $/$ 31 |

```
if RA = 0 then b }\leftarrow
else b
EA \leftarrow b + (RB)
MEM(EA, 8) \leftarrow(FRS)
```

Let the effective address (EA) be the sum (RA|0)+(RB).
The contents of register FRS are stored into the doubleword in storage addressed by EA.

Special Registers Altered:
None

## Store Floating-Point Double with Update Indexed X-form

```
stfdux FRS,RA,RB
```

| 31 | FRS | RA | RB |  | 759 | 1 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

```
EA}\leftarrow(RA)+(RB
MEM (EA, 8) \leftarrow(FRS)
RA}\leftarrow\textrm{EA
```

Let the effective address (EA) be the sum (RA)+(RB).
The contents of register FRS are stored into the doubleword in storage addressed by EA.

EA is placed into register RA.
If $R A=0$, the instruction form is invalid.
Special Registers Altered:
None

## Version 3.0

## Store Floating-Point as Integer Word Indexed X-form

stfiwx $\quad$ FRS,RA,RB

| 31 | FRS | RA | RB |  | 983 | $/$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 1 |  |  |  |  |  |  |

```
if RA = 0 then b }\leftarrow
else b
EA \leftarrow b + (RB)
MEM(EA, 4) \leftarrow(FRS) 32:63
```

Let the effective address (EA) be the sum (RA|0)+(RB). $(F R S)_{32: 63}$ are stored, without conversion, into the word in storage addressed by EA.

If the contents of register FRS were produced, either directly or indirectly, by a Load Floating-Point Single instruction, a single-precision Arithmetic instruction, or frsp, then the value stored is undefined. (The contents of register FRS are produced directly by such an instruction if FRS is the target register for the instruction. The contents of register FRS are produced indirectly by such an instruction if FRS is the final target register of a sequence of one or more Floating-Point Move instructions, with the input to the sequence having been produced directly by such an instruction.)

## Special Registers Altered:

None

### 4.6.4 Floating-Point Load and Store Double Pair Instructions [Phased-Out]

For Ifdp[x], the doubleword-pair in storage addressed by EA is loaded into an even-odd pair of FPRs with the even-numbered FPR being loaded with the leftmost doubleword from storage and the odd-numbered FPR being loaded with the rightmost doubleword.

For stfdp $[x]$, the content of an even-odd pair of FPRs is stored into the doubleword-pair in storage addressed by EA, with the even-numbered FPR being stored into the leftmost doubleword in storage and the
odd-numbered FPR being stored into the rightmost doubleword.

## Programming Note

The instructions described in this section should not be used to access an operand in DFP Extended format when the processor is in Lit-tle-Endian mode.

## Load Floating-Point Double Pair DS-form

Ifdp $\quad$ FRTp,DS(RA)

| 57 | FRTp | RA |  | DS | 00 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 |  | 11 |  | 16 |

```
if RA = 0 then b & < 
else b b (RA)
EA \leftarrow b + EXTS(DS|Ob00)
FRTP even }\leftarrow\operatorname{MEM}(EA,8
FRTP
```

Let the effective address (EA) be the sum (RA|O) + (DS\|Ob00).

The doubleword in storage addressed by EA is placed into the even-numbered register of FRTp.
The doubleword in storage addressed by EA+8 is placed into the odd-numbered register of FRTp.

If FRTp is odd, the instruction form is invalid.

## Special Registers Altered:

None

## Load Floating-Point Double Pair Indexed X-form

Ifdpx FRTp,RA,RB

| 31 | FRTp | RA | RB |  | 791 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 16 |  |  |  |  |  |

$$
\begin{aligned}
& \text { if } R A=0 \text { then } b \leftarrow 0 \\
& \text { else } b \leftarrow(R A) \\
& E A \leftarrow b+(R B) \\
& F R T p_{\text {even }} \leftarrow M E M(E A, 8) \\
& F R T p_{\text {odd }} \leftarrow \operatorname{MEM}(E A+8,8)
\end{aligned}
$$

Let the effective address (EA) be the sum (RA|0) + (RB).
The doubleword in storage addressed by EA is placed into the even-numbered register of FRTp.

The doubleword in storage addressed by EA+8 is placed into the odd-numbered register of FRTp.
If FRTp is odd, the instruction form is invalid.

## Special Registers Altered:

None

## Store Floating-Point Double Pair DS-form

stfdp $\quad$ FRSp,DS(RA)

| 61 | FRSp | RA |  | DS | 00 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 3031 |  |  |  |  |  |

```
if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(\mathrm{ RA )
EA \leftarrow b + EXTS(DS| | 0b00)
MEM(EA, 8) \leftarrowFFRSp even
MEM(EA+8, 8) \leftarrow FRSp odd
```

Let the effective address (EA) be the sum (RA|O) + (DS||Ob00).

The contents of the even-numbered register of FRSp are stored into the doubleword in storage addressed by EA.

The contents of the odd-numbered register of FRSp are stored into the doubleword in storage addressed by EA+8.

If FRSp is odd, the instruction form is invalid.

## Special Registers Altered:

None

## Store Floating-Point Double Pair Indexed X-form

stfdpx FRSp,RA,RB

| 31 | FRSp | RA | RB |  | 919 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 11 |  |  |  |  |

```
if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(\mathrm{ RA)
EA \leftarrow b + (RB)
MEM(EA, 8) \leftarrowFFRSp even
MEM(EA+8, 8) \leftarrow FRSpoda
```

Let the effective address (EA) be the sum (RA|O) + (DS||Ob00).

The contents of the even-numbered register of FRSp are stored into the doubleword in storage addressed by EA.

The contents of the odd-numbered register of FRSp are stored into the doubleword in storage addressed by EA+8.

If FRSp is odd, the instruction form is invalid.

## Special Registers Altered:

None

### 4.6.5 Floating-Point Move Instructions

These instructions copy data from one floating-point register to another, altering the sign bit (bit 0) as described below for fneg, fabs, fnabs, and fcpsgn. These instructions treat NaNs just like any other kind of
value (e.g., the sign bit of a NaN may be altered by fneg, fabs, fnabs, and fcpsgn). These instructions do not alter the FPSCR.

## Floating Negate $X$-form



The contents of register FRB with bit 0 inverted are placed into register FRT.
Special Registers Altered:
CR1
(if $\mathrm{Rc}=1$ )

## Floating Copy Sign X-form

| fabs | FRT,FRB FRT,FRB |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| fabs. |  |  |  |  |  |
| $63$ | ${ }_{6} \text { FRT }$ | $\left.\right\|_{11} \quad I / I$ | ${ }_{16}{ }^{\text {FRB }}$ | 264 | Rc <br> 31 |

The contents of register FRB with bit 0 set to zero are placed into register FRT.

Special Registers Altered:
CR1
(if $\mathrm{Rc}=1$ )

## Floating Negative Absolute Value X-form



The contents of register FRB with bit 0 set to one are placed into register FRT.

Special Registers Altered:
CR1
(if $\mathrm{Rc}=1$ )


The contents of register FRB with bit 0 set to the value of bit 0 of register FRA are placed into register FRT.
Special Registers Altered:
CR1
(if $\mathrm{Rc}=1$ )

Floating Merge Even Word X-form
fmrgew FRT,FRA,FRB

| 63 | FRT | FRA | FRB | 966 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  | 11 |  |

> if MSR.FP=O then FP_Unavailablel)
> FPR[FRT], word[0] $\leftarrow$ FPR[FRA], word[0]
> FPR[FRT]. word[1] $\leftarrow$ FPR[FRB]. word[0]

The contents of word element 0 of $\operatorname{FPR}[F R A]$ are placed into word element 0 of $\mathrm{FPR}[F R T]$.

The contents of word element 0 of $F P R[F R B]$ are placed into word element 1 of $\mathrm{FPR}[\mathrm{FRT}]$.
fmrgew is treated as a Floating-Point instruction in terms of resource availability.

## Special Registers Altered

None

## Version 3.0

## Floating Merge Odd Word X-form

fmrgow FRT,FRA,FRB

| 63 | FRT | FRA | FRB |  | 838 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |
| 31 |  |  |  |  |  |

> if MSR.FP=0 then FP_Unavailablel)
> FPR[FRT], word[0] $\leftarrow$ FPR[FRA], word[1]
> FPR[FRT]. word[1] $\leftarrow$ FPR[FRB]. wor d[1]

The contents of word element 1 of $\operatorname{FPR[FRA]}$ are placed into word element 0 of $F$ PR[FRT].

The contents of word element 1 of $\operatorname{FPR[FRB]\text {areplaced}}$ into word element 1 of $\mathrm{FPR}[F R T]$.
fmrgow is treated as a Floating-Point instruction in terms of resource availability.

## Special Registers Altered None

### 4.6.6 Floating-Point Arithmetic Instructions

### 4.6.6.1 Floating-Point Elementary Arithmetic Instructions

Floating Add [Single] A-form

| fadd | FRT,FRA,FRB FRT,FRA,FRB |  |  |  | (Rc=0) |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| fadd. |  |  |  |  |  | =1) |
| 63 | FRT | FRA | FRB | III | 21 | Rc |
| 0 | 6 | 11 | 16 | 21 | 26 | 31 |


| fadds | FRT,FRA,FRB | $(R \mathrm{c}=0)$ |
| :--- | :--- | :--- |
| fadds. | FRT,FRA,FRB | $(R \mathrm{c}=1)$ |


| 59 | FRT | FRA | FRB | I/I | 21 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 | 16 | 21 | 26 |

The floating-point operand in register FRA is added to the floating-point operand in register FRB.

If the most significant bit of the resultant significand is not 1 , the result is normalized. The result is rounded to the target precision under control of the Floating-Point Rounding Control field RN of the FPSCR and placed into register FRT.
Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( $G, R$, and $X$ ) enter into the computation.
If a carry occurs, the sum's significand is shifted right one bit position and the exponent is increased by one.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when FPSCR ${ }_{V E}=1$.

```
Special Registers Altered:
    FPRF FR FI
    FX OX UX XX
    VXSNAN VXISI
    CR1 (if Rc=1)
```

Floating Subtract [Single] A-form

| fsub | FRT,FRA,FRB | $(\mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| fsub. | FRT,FRA,FRB | $(\mathrm{Rc}=1)$ |


| 63 | FRT | FRA | FRB | I/I | 20 | Rc |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 | 26 |


| fsubs fsubs. | FRT,FRA,FRB FRT,FRA,FRB |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 59 | FRT | FRA | FRB | I/I | 20 | Rc |
| 0 | 6 | 11 | 16 | 21 | 26 | 31 |

The floating-point operand in register FRB is subtracted from the floating-point operand in register FRA.
If the most significant bit of the resultant significand is not 1 , the result is normalized. The result is rounded to the target precision under control of the Floating-Point Rounding Control field RN of the FPSCR and placed into register FRT.

The execution of the Floating Subtract instruction is identical to that of Floating Add, except that the contents of FRB participate in the operation with the sign bit (bit 0) inverted.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when $\mathrm{FPSCR}_{\mathrm{VE}}=1$.

```
Special Registers Altered:
    FPRF FR FI
    FX OX UX XX
    VXSNAN VXISI
    CR1
```

    (if \(\mathrm{Rc}=1\) )
    
## Floating Multiply [Single] A-form

| fmul fmul. | FRT,FRA,FRC FRT,FRA,FRC |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 063 | ${ }_{6}$ FRT | ${ }_{11}{ }^{\text {FRA }}$ | 16 I/I | $\begin{array}{\|c} \text { FRC } \\ 21 \end{array}$ | ${ }_{26} 25$ | $R c$ <br> 31 |



The floating-point operand in register FRA is multiplied by the floating-point operand in register FRC.

If the most significant bit of the resultant significand is not 1 , the result is normalized. The result is rounded to the target precision under control of the Floating-Point Rounding Control field RN of the FPSCR and placed into register FRT.
Floating-point multiplication is based on exponent addition and multiplication of the significands.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when $F^{F P S C R}{ }_{V E}=1$.

## Special Registers Altered:

FPRF FR FI
FX OX UX XX
VXSNAN VXIMZ
CR1
(if $\mathrm{Rc}=1$ )

## Floating Divide [Single] A-form

| fdiv fdiv. | FRT,FRA,FRB FRT,FRA,FRB |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll} 63 \\ 0 & 63 \\ \hline \end{array}$ | $6_{6}$ FRT | ${ }_{11}$ FRA | ${ }_{16}{ }^{\text {FRB }}$ | $21 / 1$ | $2^{18}$ | Rc <br> 31 |
| fdivs fdivs. | FRT,FRA,FRB FRT,FRA,FRB |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| $0 \quad 59$ | $6_{6} \text { FRT }$ | ${ }_{11}$ FRA | ${ }_{16}{ }^{\text {FRB }}$ | $21 / 1$ | ${ }^{18}$ | Rc <br> 31 |

The floating-point operand in register FRA is divided by the floating-point operand in register FRB. The remainder is not supplied as a result.

If the most significant bit of the resultant significand is not 1 , the result is normalized. The result is rounded to the target precision under control of the Floating-Point Rounding Control field RN of the FPSCR and placed into register FRT.
Floating-point division is based on exponent subtraction and division of the significands.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when $F_{P S C R}^{V E}=1$ and Zero Divide Exceptions when $\mathrm{FPSCR}_{\mathrm{ZE}}=1$.

## Special Registers Altered:

FPRF FR FI
FX OX UX ZX XX
VXSNAN VXIDI VXZDZ
CR1
(if $R c=1$ )

## Floating Square Root [Single] A-form

| fsqrt fsqrt. | FRT,FRB <br> FRT,FRB |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 63 | FRT | //] | FRB | I/I | 22 | Rc |
| 0 | 6 | 11 | 16 | 21 | 26 | 31 |


| fsqrts fsqrts. | FRT,FRB FRT,FRB |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $59$ | ${ }_{6} \text { FRT }$ | ${ }_{11} \quad / /$ | ${ }_{16}{ }^{\text {FRB }}$ | $21^{\prime \prime \prime}$ | ${ }_{26} 22$ | Rc 31 |

The square root of the floating-point operand in register FRB is placed into register FRT.

If the most significant bit of the resultant significand is not 1 , the result is normalized. The result is rounded to the target precision under control of the Floating-Point Rounding Control field RN of the FPSCR and placed into register FRT.

Operation with various special values of the operand is summarized below.

| Operand | Result | Exception |
| :---: | :---: | :---: |
| - | QNaN ${ }^{1}$ | VXSQRT |
| <0 | QNaN ${ }^{1}$ | VXSQRT |
| -0 | -0 | None |
| $+\infty$ | $+\infty$ | None |
| SNaN | QNaN ${ }^{1}$ | VXSNAN |
| QNaN | QNaN | None |
| 1 No resut | if FPSCR |  |

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when $\mathrm{FPSCR}_{\mathrm{VE}}=1$.

## Special Registers Altered:

FPRF FR FI FX OX UX XX VXSNAN VXSQRT
CR1
(if $\mathrm{Rc}=1$ )

## Floating Reciprocal Estimate [Single] A-form

| fre | FRT,FRB | $(R \mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| fre. | FRT,FRB | $(R \mathrm{c}=1)$ |


| 63 | FRT | III | FRB | I/I | 24 | Rc |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 | 26 | 31 |


| fres | FRT,FRB FRT,FRB |  |  |  | (Rc=0) |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| fres. |  |  |  |  |  | =1) |
| 59 | FRT | I/I | FRB | III | 24 | Rc |
| 0 | 6 | 11 | 16 |  | 26 | 31 |

An estimate of the reciprocal of the floating-point operand in register FRB is placed into register FRT. Unless the reciprocal would be a zero, an infinity, the result of a trap-disabled Overflow exception, or a QNaN , the estimate is correct to a precision of one part in 256 of the reciprocal of (FRB), i.e.,

$$
\operatorname{ABS}\left(\frac{\text { estimate }-1 / x}{1 / x}\right) \leq \frac{1}{256}
$$

where $x$ is the initial value in FRB.
Operation with various special values of the operand is summarized below.

| Operand | Result | Exception |
| :---: | :---: | :---: |
| - - | -0 | None |
| -0 | $-\infty^{1}$ | ZX |
| +0 | $+\infty^{1}$ | ZX |
| $+\infty$ | +0 | None |
| SNaN | QNaN ${ }^{2}$ | VXSNAN |
| QNaN | QNaN | None |
| 1 No result if $\mathrm{FPSCR}_{\text {ZE }}=1$. |  |  |
| 2 No result | FPSCR |  |

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when FPSCR $_{V E}=1$ and Zero Divide Exceptions when $\mathrm{FPSCR}_{\mathrm{ZE}}=1$.

The results of executing this instruction may vary between implementations, and between different executions on the same implementation.

## Special Registers Altered:

FPRF FR (undefined) FI (undefined)
FX OX UX ZX XX (undefined)
VXSNAN
CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

For the Floating-Point Estimate instructions, some implementations might implement a precision higher than the minimum architected precision. Thus, a program may take advantage of the higher precision instructions to increase performance by decreasing the iterations needed for software emulation of floating-point instructions. However, there is no guarantee given about the precision which may vary (up or down) between implementations. Only programs targeted at a specific implementation (i.e., the program will not be migrated to another implementation) should take advantage of the higher precision of the instructions. All other programs should rely on the minimum architected precision, which will guarantee the program to run properly across different implementations.

## Floating Reciprocal Square Root Estimate [Single] A-form

| frsqrte | FRT,FRB | (Rc=0) |
| :--- | :--- | :--- |
| frsqrte. | FRT,FRB | (Rc=1) |


| ${ }_{0} 63$ | ${ }_{6} \text { FRT }$ | $11 /$ | $\begin{aligned} & \text { FRB } \\ & 16 \end{aligned}$ |  | 26 | Rc |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |


| frsqrtes | FRT,FRB | (Rc=0) |
| :--- | :--- | :--- |
| frsqrtes. | FRT,FRB | $(R c=1)$ |


| 59 | ${ }_{6}$ FRT | 11 | 16 | 21 | ${ }_{26}{ }^{26}$ | Rc 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |

A estimate of the reciprocal of the square root of the floating-point operand in register FRB is placed into register FRT. The estimate placed into register FRT is correct to a precision of one part in 32 of the reciprocal of the square root of (FRB), i.e.,

$$
\operatorname{ABS}\left(\frac{\text { estimate }-1 /(\sqrt{x})}{1 /(\sqrt{x})}\right) \leq \frac{1}{32}
$$

where $x$ is the initial value in FRB.
Operation with various special values of the operand is summarized below.

| Operand | Result | Exception |
| :---: | :---: | :---: |
| - | QNaN ${ }^{2}$ | VXSQRT |
| $<0$ | QNaN ${ }^{2}$ | VXSQRT |
| -0 | $-\infty^{1}$ | ZX |
| +0 | $+\infty^{1}$ | ZX |
| $+\infty$ | +0 | None |
| SNaN | QNaN ${ }^{2}$ | VXSNAN |
| QNaN | QNaN | None |
| 1 No result if FPSCR ${ }_{\text {ZE }}=1$. |  |  |
| 2 No result | if FPSCR |  |

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when FPSCR $_{\text {VE }}=1$ and Zero Divide Exceptions when $\mathrm{FPSCR}_{\mathrm{ZE}}=1$.
The results of executing this instruction may vary between implementations, and between different executions on the same implementation.

## Special Registers Altered:

FPRF FR (undefined) FI (undefined)
FX OX UX ZX XX (undefined)
VXSNAN VXSQRT
CR1
(if $R c=1$ )

## Note

See the Notes that appear with fre[s].

## Floating Test for software Divide X-form

ftdiv
BF,FRA,FRB

| 63 | BF | $/ 1$ | FRA | FRB |  | 128 |
| :---: | :---: | :--- | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 9 | 11 |  |  |

Let e_a be the unbiased exponent of the double-precision floating-point operand in register FRA.

Let e_b be the unbiased exponent of the double-precision floating-point operand in register FRB.
fe_flag is set to 1 if any of the following conditions occurs.

■ The double-precision floating-point operand in register FRA is a NaN or an Infinity.

■ The double-precision floating-point operand in register FRB is a Zero, a NaN, or an Infinity.

- e_b is less than or equal to -1022.

■ e_b is greater than or equal to 1021.
■ The double-precision floating-point operand in register FRA is not a zero and the difference, e_a - e_b, is greater than or equal to 1023.

■ The double-precision floating-point operand in register FRA is not a zero and the difference, e_a - e_b, is less than or equal to -1021.

- The double-precision floating-point operand in register FRA is not a zero and e_a is less than or equal to -970

Otherwise fe_flag is set to 0 .
$f g \_f l a g$ is set to 1 if either of the following conditions occurs.

■ The double-precision floating-point operand in register FRA is an Infinity.

- The double-precision floating-point operand in register FRB is a Zero, an Infinity, or a denormalized value.
Otherwise fg_flag is set to 0 .
If the implementation guarantees a relative error of fre[s][.] of less than or equal to $2^{-14}$, then fl_flag is set to 1 . Otherwise fl_flag is set to 0 .
$C R$ field $B F$ is set to the value fl_flag || fg_flag || fe_flag || 0 b0.


## Special Registers Altered:

CR field $B F$

## Floating Test for software Square Root X-form

ftsqrt $\quad B F, F R B$

| 63 | BF | // |  | //I | FRB |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 9 |  | 6 | 160 | $1 /$ |  |  |

Let e_b be the unbiased exponent of the double-precision floating-point operand in register FRB.
fe_flag is set to 1 if either of the following conditions occurs.

- The double-precision floating-point operand in register FRB is a zero, a NaN, or an infinity, or a negative value.

■ e_b is less than or equal to -970.
Otherwise fe_flag is set to 0 .
fg_flag is set to 1 if the following condition occurs.

- The double-precision floating-point operand in register FRB is a Zero, an Infinity, or a denormalized value.

Otherwise fg_flag is set to 0 .
If the implementation guarantees a relative error of frsqrte[s][.] of less than or equal to $2^{-14}$, then fl_flag is set to 1 . Otherwise fl_flag is set to 0 .
$C R$ field $B F$ is set to the value fl_flag || fg_flag || fe_flag || 0b0.

## Special Registers Altered:

CR field BF

## Programming Note

ftdiv and ftsqrt are provided to accelerate software emulation of divide and square root operations, by performing the requisite special case checking. Software needs only a single branch, on $\mathrm{FE}=1$ (in $\mathrm{CR}[B F])$, to a special case handler. FG and FL may provide further acceleration opportunities.

### 4.6.6.2 Floating-Point Multiply-Add Instructions

These instructions combine a multiply and an add operation without an intermediate rounding operation. The fraction part of the intermediate product is 106 bits wide (L bit, FRACTION), and all 106 bits take part in the add/ subtract portion of the instruction.

Status bits are set as follows.
■ Overflow, Underflow, and Inexact Exception bits, the FR and FI bits, and the FPRF field are set
based on the final result of the operation, and not on the result of the multiplication.
■ Invalid Operation Exception bits are set as if the multiplication and the addition were performed using two separate instructions (fmu[s], followed by $\boldsymbol{f a d d}[\boldsymbol{s}]$ or $\boldsymbol{f s u b}[\boldsymbol{s}]$ ). That is, multiplication of infinity by 0 or of anything by an SNaN , and/or addition of an SNaN , cause the corresponding exception bits to be set.

## Floating Multiply-Add [Single] A-form

| fmadd | FRT,FRA,FRC,FRB | $(R \mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| fmadd. | FRT,FRA,FRC,FRB | $(R c=1)$ |


| 63 | FRT | FRA | FRB | FRC | 29 | Rc |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 | 26 | 31 |


| fmadds | FRT,FRA,FRC,FRB | $(\mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| fmadds. | FRT,FRA,FRC,FRB | $(R \mathrm{Rc}=1)$ |


| 59 | FRT | FRA | FRB | FRC | 29 | Rc |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 | 26 |
| 12 |  |  |  |  |  |  |

The operation

$$
\mathrm{FRT} \leftarrow[(\mathrm{FRA}) \times(\mathrm{FRC})]+(\mathrm{FRB})
$$

is performed.
The floating-point operand in register FRA is multiplied by the floating-point operand in register FRC. The float-ing-point operand in register FRB is added to this intermediate result.

If the most significant bit of the resultant significand is not 1 , the result is normalized. The result is rounded to the target precision under control of the Floating-Point Rounding Control field RN of the FPSCR and placed into register FRT.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when $\mathrm{FPSCR}_{\mathrm{VE}}=1$.

```
Special Registers Altered:
    FPRF FR FI
    FX OX UX XX
    VXSNAN VXISI VXIMZ
    CR1
    (if Rc=1)
```

Floating Multiply-Subtract [Single] A-form

| fmsub | FRT,FRA,FRC,FRB | (Rc=0) |
| :--- | :--- | :--- |
| fmsub. | FRT,FRA,FRC,FRB | $(R c=1)$ |


| 63 | FRT | FRA | FRB | FRC | 28 | RC |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 | 26 |


| fmsubs | FRT,FRA,FRC,FRB | $(\mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| fmsubs. | FRT,FRA,FRC,FRB | $(\mathrm{Rc}=1)$ |


| 59 | FRT | FRA | FRB | FRC | 28 | RC |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 | 26 | 31 |

The operation

$$
\mathrm{FRT} \leftarrow[(\mathrm{FRA}) \times(\mathrm{FRC})]-(\mathrm{FRB})
$$

is performed.
The floating-point operand in register FRA is multiplied by the floating-point operand in register FRC. The float-ing-point operand in register FRB is subtracted from this intermediate result.

If the most significant bit of the resultant significand is not 1 , the result is normalized. The result is rounded to the target precision under control of the Floating-Point Rounding Control field RN of the FPSCR and placed into register FRT.
FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when $\mathrm{FPSCR}_{\mathrm{VE}}=1$.

## Special Registers Altered:

FPRF FR FI
FX OX UX XX
VXSNAN VXISI VXIMZ
CR1
(if $\mathrm{Rc}=1$ )

## Floating Negative Multiply-Add [Single] A-form

| fnmadd | FRT,FRA,FRC,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| fnmadd. | FRT,FRA,FRC,FRB | $(R c=1)$ |


| 63 | FRT | FRA | FRB | FRC | 31 | Rc |
| :--- | :--- | :--- | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 | 26 | 31 |


| fnmadds | FRT,FRA,FRC,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| fnmadds. | FRT,FRA,FRC,FRB | $(R c=1)$ |


| 59 | FRT | FRA | FRB | FRC | 31 | Rc |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 | 26 |

The operation

$$
\text { FRT } \leftarrow-([(F R A) \times(F R C)]+(F R B))
$$

is performed.
The floating-point operand in register FRA is multiplied by the floating-point operand in register FRC. The float-ing-point operand in register FRB is added to this intermediate result.
If the most significant bit of the resultant significand is not 1 , the result is normalized. The result is rounded to the target precision under control of the Floating-Point Rounding Control field RN of the FPSCR, then negated and placed into register FRT.

This instruction produces the same result as would be obtained by using the Floating Multiply-Add instruction and then negating the result, with the following exceptions.

■ QNaNs propagate with no effect on their "sign" bit.

- QNaNs that are generated as the result of a disabled Invalid Operation Exception have a "sign" bit of 0 .
- SNaNs that are converted to QNaNs as the result of a disabled Invalid Operation Exception retain the "sign" bit of the SNaN.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when $\mathrm{FPSCR}_{\mathrm{VE}}=1$.

## Special Registers Altered:

FPRF FR FI
FX OX UX XX
VXSNAN VXISI VXIMZ
CR1
(if $\mathrm{Rc}=1$ )

## Floating Negative Multiply-Subtract [Single] A-form

| fnmsub | FRT,FRA,FRC,FRB | $(\mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| fnmsub. | FRT,FRA,FRC,FRB | $(\mathrm{Rc}=1)$ |


| 63 | FRT | FRA | FRB | FRC | 30 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 | 16 | 21 | 26 |
| 31 |  |  |  |  |  |  |


| fnmsubs | FRT,FRA,FRC,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| fnmsubs. | FRT,FRA,FRC,FRB | $(R c=1)$ |


| 59 | FRT | FRA | FRB | FRC | 30 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 | 16 | 21 | 26 |
| 1 |  |  |  |  |  |  |

The operation

$$
\text { FRT } \leftarrow-([(F R A) \times(F R C)]-(F R B))
$$

is performed.
The floating-point operand in register FRA is multiplied by the floating-point operand in register FRC. The float-ing-point operand in register FRB is subtracted from this intermediate result.

If the most significant bit of the resultant significand is not 1 , the result is normalized. The result is rounded to the target precision under control of the Floating-Point Rounding Control field RN of the FPSCR, then negated and placed into register FRT.

This instruction produces the same result as would be obtained by using the Floating Multiply-Subtract instruction and then negating the result, with the following exceptions.

■ QNaNs propagate with no effect on their "sign" bit.

- QNaNs that are generated as the result of a disabled Invalid Operation Exception have a "sign" bit of 0 .
- SNaNs that are converted to QNaNs as the result of a disabled Invalid Operation Exception retain the "sign" bit of the SNaN.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when $\mathrm{FPSCR}_{\mathrm{VE}}=1$.

## Special Registers Altered:

FPRF FR FI
FX OX UX XX
VXSNAN VXISI VXIMZ
CR1
(if $\mathrm{Rc}=1$ )

### 4.6.7 Floating-Point Rounding and Conversion Instructions

## Programming Note

Examples of uses of these instructions to perform various conversions can be found in Section E.2, "Floating-Point Conversions" on page 642.

### 4.6.7.1 Floating-Point Rounding Instruction

## Floating Round to Single-Precision X-form



The floating-point operand in register $F R B$ is rounded to single-precision, using the rounding mode specified by RN, and placed into register FRT.

The rounding is described fully in Section A.1, "Float-ing-Point Round to Single-Precision Model" on page 779.

FPRF is set to the class and sign of the result, except for Invalid Operation Exceptions when $V E=1$.
Special Registers Altered:

```
FPRF FR FI
    FX OX UX XX VXSNAN
    CR1
```


### 4.6.7.2 Floating-Point Convert To/From Integer Instructions

## Floating Convert To Integer Doubleword X-form



Let src be the double-precision floating-point value in FRB.

If src is a NaN , then the result is $0 \times 8000-0000 \_0000 \_0000, \mathrm{VXCVI}$ is set to 1 , and, if src is an SNäN, VX'SNAN is set to 1.
Otherwise, src is rounded to a floating-point integer using the rounding mode specified by RN.
If the rounded value is greater than $2^{63} \cdot 1$, then the result is $0 \times 7$ FFF_ FFFF_FFF_FFFF and VXCVI is set to 1 .
Otherwise, if the rounded value is less than $\cdot 2^{63}$, then the result is $0 \times 8000 \_0000 \_0000 \_0000$ and VXCVI is set to 1.

Otherwise, the result is the rounded value converted to 64 -bit signed-integer format, and $X X$ is set to 1 if the result is inexact.

If an enabled Invalid Operation Exception does not occur, then the result is placed into $F$ RT.

The conversion is described fully in Section A.2, "Float-ing-Point Convert to Integer Model" on page 783.
Except for enabled Invalid Operation Exceptions, FPRF is undefined. $F R$ is set if the result is incremented when rounded. Fl is set if the result is inexact.

## Special Registers Altered:

FPRF (undefined) FR FI
FX XX VXSNAN VXCVI
CR1 (if $R C=1$ )

## Floating Convert To Integer Doubleword with round toward Zero X-form

```
fctidz FRT,FRB (RC=0)
fctidz. FRT,FRB (RC=1)
```

| 63 | FRT |  | FRB |  | 815 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 11 | 16 | 21 |  |

Let src be the double-precision floating-point value in FRB.

If src is a NaN , then the result is $0 \times 8000000000000000$, VXCVI is set to 1 , and, if src is an SNaN, VXSNAN is set to 1 .

Otherwise, src is rounded to a floating-point integer using the rounding mode Round toward Zero
If the rounded value is greater than $2^{63} .1$, then the result is $0 \times 7$ FFF_FFFF_FFF_FFFF and VXCVI is set to 1 .
Otherwise, if the rounded value is less than $\cdot 2^{63}$, then the result is $0 \times 80000_{-} 00000_{-} 00000_{-} 0000$ and VXCVI is set to 1.

Otherwise, the result is the rounded value converted to 64 -bit signed-integer format, and $X X$ is set to 1 if the result is inexact.

If an enabled Invalid Operation Exception does not occur, then the result is placed into FRT.
The conversion is described fully in Section A.2, "Float-ing-Point Convert to Integer Model" on page 783.

Except for enabled Invalid Operation Exceptions, FPRF is undefined. $F R$ is set if the result is incremented when rounded. Fl is set if the result is inexact.

## Special Registers Altered:

$$
\begin{aligned}
& \text { FPRF (undefined) FR FI } \\
& \text { FX XX VXSNAN VXCVI } \\
& \text { CR1 }
\end{aligned}
$$

(if $R C=1$ )

## Floating Convert To Integer Doubleword Unsigned X-form

| fctidu | FRT,FRB | $(R C=0)$ |
| :--- | :--- | :--- |
| fctidu. | FRT,FRB | $(R C=1)$ |


| 63 | FRT |  | I/I | FRB |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 642 | Rc |  |  |
| 31 |  |  |  |  |  |

Let src be the double-precision floating-point value in FRB.

If $\operatorname{src}$ is a NaN , then the result is $0 \times 0000 \_0000 \_0000 \_0000$, VXCVI is set to 1 , and, if src is an SNāN, VX̄SNAN is set to 1 .

Otherwise, src is rounded to a floating-point integer using the rounding mode specified by RN.
If the rounded value is greater than $2^{64} .1$, then the result is $0 \times$ FFFF_FFFF_FFF_FFFF, and VXCVI is set to 1 .

Otherwise, if the rounded value is less than 0 , then the result is $0 \times 00000_{-} 0000_{-} 0000 \_0000$, and VXCVI is set to 1 .

Otherwise, the result is the rounded value converted to 64-bit unsigned-integer format, and $X X$ is set to 1 if the result is inexact.

If an enabled Invalid Operation Exception does not occur, then the result is placed into FRT.

The conversion is described fully in Section A.2, "Float-ing-Point Convert to Integer Model" on page 783.

Except for enabled Invalid Operation Exceptions, FPRF is undefined. $F R$ is set if the result is incremented when rounded. Fl is set if the result is inexact.

## Special Registers Altered:

FPRF (undefined) FR FI
FX XX VXSNAN VXCVI
CR1
(if $R c=1$ )

## Floating Convert To Integer Doubleword Unsigned with round toward Zero X-form

## Floating Convert To Integer Word X-form

| fctiw fctiw. | FRT,FRB FRT,FRB |  |  |  | $\begin{aligned} & (R C=0) \\ & (R C=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 63 | FRT | I/I | FRB |  | 14 | Rc |
| 0 | 6 | 11 | 16 | 21 |  | 31 |

Let $\operatorname{src}$ be the double-precision floating-point value in FRB.

If $5 r \mathrm{C}$ is a NaN , then the result is $0 \times 8000 \_0000, \mathrm{VXCVI}$ is set to 1 , and, if src is an SNaN, VXSNAN is set to 1.

Otherwise, src is rounded to a floating-point integer using the rounding mode specified by RN.
If the rounded value is greater than $2^{31} .1$, then the result is $0 \times 7$ FFF_FFFF, and VXCVI is set to 1 .
Otherwise, if the rounded value is less than $\cdot 2^{31}$, then the result is $0 \times 8000 \_0000$, and VXCVI is set to 1 .
Otherwise, the result is the rounded value converted to 32 -bit signed-integer format, and $X X$ is set to 1 if the result is inexact.

If an enabled Invalid Operation Exception does not occur, then the result is placed into $F R T_{32: 63}$ and $F R T_{0: 31}$ is undefined,

The conversion is described fully in Section A.2, "Float-ing-Point Convert to Integer Model" on page 783.
Except for enabled Invalid Operation Exceptions, FPRF is undefined. $F R$ is set if the result is incremented when rounded. Fl is set if the result is inexact.

## Special Registers Altered:

FPRF (undefined) FR Fl
FX XX VXSNAN VXCVI
CR1
(if $R C=1$ )

## Floating Convert To Integer Word with round toward Zero X-form

fctiwz
FRT,FRB
( $\mathrm{Rc}=0$ )
fctiwz. $\operatorname{FRT}, \mathrm{FRB}(R c=1)$ Let src be the double-precision

| 63 | FRT | ${ }^{\text {I }}$ //I | FRB |  | 15 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

floating-point value in $F$ RB.
If src is a NaN , then the result is $0 \times 80000_{-} 0000, \mathrm{VXCVI}$ is set to 1 , and, if src is an $\mathrm{SNaN}, \mathrm{VXSNAN}$ is set to 1 .

Otherwise, $\operatorname{src}$ is rounded to a floating-point integer using the rounding mode Round toward Zero.
If the rounded value is greater than $2^{31} \cdot 1$, then the result is $0 \times 7$ FFF_FFFF, and VXCVI is set to 1 .
Otherwise, if the rounded value is less than $\cdot 2^{31}$, then the result is $0 \times 8000.0000$, and VXCVI is set to 1 .
Otherwise, the result is the rounded value converted to 32 -bit signed-integer format, and $X X$ is set to 1 if the result is inexact.

If an enabled Invalid Operation Exception does not occur, then the result is placed into $F R T_{32: 63}$ and $F R T_{0: 31}$ is undefined,
The conversion is described fully in Section A.2, "Float-ing-Point Convert to Integer Model" on page 783.

Except for enabled Invalid Operation Exceptions, FPRF is undefined. $F R$ is set if the result is incremented when rounded. Fl is set if the result is inexact.

## Special Registers Altered:

$\begin{array}{lll}\text { FPRF (undefined) FR FI } & \\ \text { FX XX } & \\ \text { VXSNAN VXCVI } & & \\ \text { CR1 } & \text { (if } R C=1)\end{array}$

## Floating Convert To Integer Word Unsigned $X$-form

| fctiwu | FRT,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| fctiwu. | FRT,FRB | $(R c=1)$ |


| 63 | FRT |  | FRB |  | 142 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |
| 10 |  |  |  |  |  |

Let $\operatorname{src}$ be the double-precision floating-point value in FRB.

If $5 r_{\mathrm{C}}$ is a NaN , then the result is $0 \times 00000_{0} 0000$, VXCVI is set to 1 , and, if $s r^{\circ}$ is an SNaN, VXSNAN is set to 1 .

Otherwise, src is rounded to a floating-point integer using the rounding mode specified by RN.
If the rounded value is greater than $2^{32} \cdot 1$, then the result is $0 \times F F F F$ FFFF and VXCVI is set to 1 .

Otherwise, if the rounded value is less than 0 , then the result is $0 \times 0000 \_0000$ and VXCVI is set to 1 .
Otherwise, the result is the rounded value converted to 32-bit unsigned-integer format, and $X X$ is set to 1 if the result is inexact.

If an enabled Invalid Operation Exception does not occur, then the result is placed into $F R T_{32: 63}$ and $F R T_{0: 31}$ is undefined,

The conversion is described fully in Section A.2, "Float-ing-Point Convert to Integer Model" on page 783.
Except for enabled Invalid Operation Exceptions, FPRF is undefined. $F R$ is set if the result is incremented when rounded. Fl is set if the result is inexact.

## Special Registers Altered:

FPRF (undefined) FR FI
FX XX
VXSNAN VXCVI
CR1

## Floating Convert To Integer Word Unsigned with round toward Zero X-form

| fctiwuz fctiwuz. | FRT,FRB FRT,FRB |  |  | $\begin{aligned} & (R C=0) \\ & (R C=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $063$ | $6$ | $11 /$ | ${ }_{16}$ FRB | $\begin{array}{ll} 21 & 143 \\ \hline 21 \end{array}$ | Rc <br> 31 |

Let $\operatorname{src}$ be the double-precision floating-point value in FRB.

If 5 Cr is a NaN , then the result is $0 \times 0000 \_0000$, VXCVI is set to 1 , and, if $\operatorname{src}$ is an SNaN, VXSNAN is set to 1 .

Otherwise, $\operatorname{src}$ is rounded to a floating-point integer using the rounding mode Round toward Zero.
If the rounded value is greater than $2^{32} \cdot 1$, then the result is $0 \times F F F F$ FFFF and $V X C V I$ is set to 1 .

Otherwise, if the rounded value is less than 0.0 , then the result is $0 \times 0000 \_0000$ and VXCVI is set to 1 .
Otherwise, the result is the rounded value converted to 32 -bit unsigned-integer format, and $X X$ is set to 1 if the result is inexact.

If an enabled Invalid Operation Exception does not occur, then the result is placed into $F R T_{32: 63}$ and $F R T_{0: 31}$ is undefined,

The conversion is described fully in Section A.2, "Float-ing-Point Convert to Integer Model" on page 783.
Except for enabled Invalid Operation Exceptions, FPRF is undefined. $F R$ is set if the result is incremented when rounded. Fl is set if the result is inexact.

## Special Registers Altered:

$\begin{array}{ll}\text { FPRF (undefined) FR FI } \\ \text { FX XX VXSNAN VXCVI } & \\ \text { CR1 } & \text { (if } R C=1)\end{array}$

## Floating Convert From Integer Doubleword X-form

| fcfid <br> fcfid.$\quad$FRT,FRB <br> FRT,FRB |
| :--- |
| 63 FRT I/I FRB  $(R C=0)$ <br> $(R c=1)$      |
| 0 |

The 64-bit signed fixed-point operand in register FRB is converted to an infinitely precise floating-point integer. The result of the conversion is rounded to double-precision, using the rounding mode specified by RN, and placed into register FRT.

The conversion is described fully in Section A.3, "Float-ing-Point Convert from Integer Model".
$F P R F$ is set to the class and sign of the result. $F R$ is set if the result is incremented when rounded. Fl is set if the result is inexact.

## Special Registers Altered:

FPRF FR FI FX XX CR1

$$
\text { (if } R C=1 \text { ) }
$$

## Programming Note

Converting a signed integer word to double-precision floating-point can be accomplished by loading the word from storage using Load Float Word Algebraic Indexed and then using fcfid.

## Floating Convert From Integer Doubleword Unsigned X-form

| fcfidu | FRT,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| fcfidu. | FRT,FRB | $(R c=1)$ |


| 63 | FRT |  | FRB |  | 974 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 11 | 16 | 21 |  |

The 64-bit unsigned fixed-point operand in register FRB is converted to an infinitely precise floating-point integer. The result of the conversion is rounded to dou-ble-precision, using the rounding mode specified by FPSCR ${ }_{R N}$, and placed into register FRT.

The conversion is described fully in Section A.3, "Float-ing-Point Convert from Integer Model".
FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result. FR is set if the result is incremented when rounded. FPSCR $_{\text {FI }}$ is set if the result is inexact.

## Special Registers Altered:

FPRF FR FI
FX XX
CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

Converting an unsigned integer word to dou-ble-precision floating-point can be accomplished by loading the word from storage using Load Float Word and Zero Indexed and then using fcfidu.

## Floating Convert From Integer Doubleword Single X-form

| fcfids | FRT,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| fcfids. | FRT,FRB | $(R c=1)$ |


| 59 | FRT | I/I | FRB | 846 | Rc |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 11 | 16 |  | 31 |

The 64-bit signed fixed-point operand in register FRB is converted to an infinitely precise floating-point integer. The result of the conversion is rounded to single-precision, using the rounding mode specified by $\mathrm{FPSCR}_{\mathrm{RN}}$, and placed into register FRT.

The conversion is described fully in Section A.3, "Float-ing-Point Convert from Integer Model".

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result. $F R$ is set if the result is incremented when rounded. FPSCR $_{\text {FI }}$ is set if the result is inexact.

```
Special Registers Altered:
        FPRF FR FI
        FX XX
        CR1
                            (if Rc=1)
```


## Programming Note

Converting a signed integer word to single-precision floating-point can be accomplished by loading the word from storage using Load Float Word Algebraic Indexed and then using fcfids.

## Floating Convert From Integer Doubleword Unsigned Single X-form

| fcfidus | FRT,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| fcfidus. | FRT,FRB | $(R c=1)$ |



The 64-bit unsigned fixed-point operand in register FRB is converted to an infinitely precise floating-point integer. The result of the conversion is rounded to sin-gle-precision, using the rounding mode specified by FPSCR ${ }_{\text {RN }}$, and placed into register FRT.

The conversion is described fully in Section A.3, "Float-ing-Point Convert from Integer Model".
FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result. FR is set if the result is incremented when rounded. FPSCR $_{\text {FI }}$ is set if the result is inexact.

## Special Registers Altered:

FPRF FR FI
FX XX
(if $\mathrm{Rc}=1$ )

### 4.6.7.3 Floating Round to Integer Instructions

The Floating Round to Integer instructions provide direct support for rounding functions found in high level languages. For example, frin, friz, frip, and frim implement C++ round(), trunc(), ceil(), and floor(), respectively. Note that frin does not implement the IEEE Round to Nearest function, which is often further described as "ties to even." The rounding performed by these instructions is described fully in Section A.4, "Floating-Point Round to Integer Model" on page 788.

## Programming Note

These instructions set FPSCR FR FI to $0 b 00$ regardless of whether the result is inexact or rounded because there is a desire to preserve the value of $\mathrm{FPSCR}_{\mathrm{XX}}$. Furthermore, it is believed that most programs do not need to know whether these rounding operations produce inexact or rounded results. If it is necessary to determine whether the result is inexact or rounded, software must compare the result with the original source operand.

## Floating Round to Integer Nearest X-form



The floating-point operand in register FRB is rounded to an integral value as follows, with the result placed into register FRT. If the sign of the operand is positive, (FRB) +0.5 is truncated to an integral value, otherwise (FRB) - 0.5 is truncated to an integral value.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when $\mathrm{FPSCR}_{\mathrm{VE}}=1$.

```
Special Registers Altered:
FPRF FR (set to 0) FI (set to 0)
FX
VXSNAN
CR1
```

(if $\mathrm{Rc}=1$ )

Floating Round to Integer Toward Zero X-form


The floating-point operand in register FRB is rounded to an integral value using the rounding mode round toward zero, and the result is placed into register FRT.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when $\mathrm{FPSCR}_{\mathrm{VE}}=1$.

```
Special Registers Altered:
    FPRF FR (set to 0) FI (set to 0)
    FX
    VXSNAN
    CR1
    (if Rc=1)
```


## Floating Round to Integer Plus X-form

| frip frip. | FRT,FRB |  |  |  | $\begin{aligned} & (R c=0) \\ & (R c=1) \end{aligned}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | FRT |  |  |  |  |
| 63 | FRT | //I | FRB | 456 | Rc |
| 0 |  |  | 6 |  | 31 |

The floating-point operand in register FRB is rounded to an integral value using the rounding mode round toward +infinity, and the result is placed into register FRT.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when FPSCR $_{\text {VE }}=1$.
Special Registers Altered:
FPRF FR (set to 0) FI (set to 0)
FX
VXSNAN
CR1
(if $\mathrm{Rc}=1$ )

Floating Round to Integer Minus X-form

| frim |
| :--- |
| frim |
| frim. | FRT,FRB

FRT,FRB $\quad$| (Rc=0) |
| :--- |
| $(R c=1)$ |

The floating-point operand in register FRB is rounded to an integral value using the rounding mode round toward -infinity, and the result is placed into register FRT.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exceptions when FPSCR $_{\text {VE }}=1$.

## Special Registers Altered:

FPRF FR (set to 0) FI (set to 0)
FX
VXSNAN
CR1

### 4.6.8 Floating-Point Compare Instructions

The floating-point Compare instructions compare the contents of two floating-point registers. Comparison ignores the sign of zero (i.e., regards +0 as equal to -0 ). The comparison can be ordered or unordered.
The comparison sets one bit in the designated CR field to 1 and the other three to 0 . The FPCC is set in the same way.

The CR field and the FPCC are set as follows.

| Bit | Name | Description |
| :--- | :--- | :--- |
| 0 | FL | (FRA) $<$ (FRB) |
| 1 | FG | (FRA) $>$ (FRB) |
| 2 | FE | (FRA) $=$ (FRB) |
| 3 | FU | (FRA) ? (FRB) (unordered) |

## Floating Compare Unordered X-form

fcmpu BF,FRA,FRB

| 63 | BF | $/ /$ | FRA | FRB |  | 0 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 9 | 9 | 11 | 16 | 21 |  | 31 |

```
if (FRA) is a NaN or
    (FRB) is a NaN then c \leftarrow0b0001
else if (FRA) < (FRB) then c \leftarrow 0b1000
else if (FRA) > (FRB) then c < 0b0100
else c\leftarrow0b0010
FPCC }\leftarrow\textrm{c
CR 4\timesBF:4\timesBF+3
if (FRA) is an SNaN or
    (FRB) is an SNaN then
        VXSNAN }\leftarrow
```

The floating-point operand in register FRA is compared to the floating-point operand in register FRB. The result of the compare is placed into CR field BF and the FPCC.

If either of the operands is a NaN , either quiet or signaling, then CR field BF and the FPCC are set to reflect unordered. If either of the operands is a Signaling NaN , then VXSNAN is set.

## Special Registers Altered:

CR field BF
FPCC
FX
VXSNAN

Floating Compare Ordered X-form


The floating-point operand in register FRA is compared to the floating-point operand in register FRB. The result of the compare is placed into CR field BF and the FPCC.

If either of the operands is a NaN , either quiet or signaling, then CR field BF and the FPCC are set to reflect unordered. If either of the operands is a Signaling NaN , then VXSNAN is set and, if Invalid Operation is disabled (VE=0), VXVC is set. If neither operand is a Signaling NaN but at least one operand is a Quiet NaN , then VXVC is set.

```
Special Registers Altered:
    CR field BF
    FPCC
    FX
    VXSNAN VXVC
```


### 4.6.9 Floating-Point Select Instruction

## Floating Select A-form

| fsel <br> fsel. | FRT,FRA,FRC,FRB FRT,FRA,FRC,FRB |  |  | $\begin{aligned} & (R c=0) \\ & (R c=1) \end{aligned}$ |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 63 | FRT | FRA | FRB | FRC | 23 | Rc |
| 0 | 6 | 11 | 16 | 21 | 26 | 31 |

```
if (FRA) \geq 0.0 then FRT \leftarrow (FRC)
else FRT \leftarrow (FRB)
```

The floating-point operand in register FRA is compared to the value zero. If the operand is greater than or equal to zero, register FRT is set to the contents of register FRC . If the operand is less than zero or is a NaN , register FRT is set to the contents of register FRB. The com-
parison ignores the sign of zero (i.e., regards +0 as equal to -0 ).

## Special Registers Altered:

CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

Examples of uses of this instruction can be found in Sections E.2, "Floating-Point Conversions" on page 642 and E.3, "Floating-Point Selection" on page 646.

Warning: Care must be taken in using fsel if IEEE compatibility is required, or if the values being tested can be NaNs or infinities; see Section E.3.4, "Notes" on page 646.

This section gives examples of how the Floating Select instruction can be used to implement certain simple forms of if-then-else constructions, without branching.

The examples show program fragments in an imaginary, C-like, high-level programming language, and the corresponding program fragment using fsel and other Power ISA instructions. In the examples, a, b, x, y, and $z$ are float-ing-point variables, which are assumed to be in FPRs $f a, f b, f x, f y$, and $f z$. FPR fs is assumed to be available for scratch space.

Warning: Care must be taken in using fsel if IEEE compatibility is required, or if the values being tested can be NaNs or infinities; see Section .

## Comparison to Zero

| High-level language: | Power ISA: | Notes |
| :---: | :---: | :---: |
| ```if a \geq 0.0 then x}\leftarrow else x}\leftarrow\textrm{z``` | fsel fx,fa,fy,fz | (1) |
| ```if a> 0.0 then x}\leftarrow else x}\leftarrow\textrm{z``` | $\begin{array}{ll} \text { fneg } & \mathrm{fs}, \mathrm{fa} \\ \text { fsel } & \mathrm{fx}, \mathrm{fs}, \mathrm{fz}, \mathrm{fy} \end{array}$ | $(1,2)$ |
| $\begin{aligned} & \text { if } a=0.0 \text { then } \\ & x \\ & x \nleftarrow y \\ & \text { else } \\ & \quad x \nleftarrow z \end{aligned}$ | $\begin{aligned} & \text { fsel } \mathrm{fx}, \mathrm{fa}, \mathrm{fy}, \mathrm{fz} \\ & \text { fneg } \mathrm{fs}, \mathrm{fa} \\ & \text { fsel } \mathrm{fx}, \mathrm{fs}, \mathrm{fx}, \mathrm{fz} \end{aligned}$ | (1) |

## Notes:

The following Notes apply to the preceding examples and to the corresponding cases using the other three arithmetic relations $(<, \leq$, and $\neq)$. They should also be considered when any other use of $\boldsymbol{f s e l}$ is contemplated.

In these Notes, the "optimized program" is the Power ISA program shown, and the "unoptimized program" (not shown) is the corresponding Power ISA program that uses fcmpu and Branch Conditional instructions instead of fsel.

## Simple if-then-else Constructions

| High-level language: | Power ISA: | Notes |
| :---: | :---: | :---: |
| $\text { if } \mathrm{a} \geq \mathrm{b} \text { then } \mathrm{x} \leftarrow \mathrm{y}$ $\text { else } x \leftarrow z$ | fsub fs, fa, fb <br> fsel fx,fs,fy,fz | $(4,5)$ |
| if $\mathrm{a}>\mathrm{b}$ then $\mathrm{x} \leftarrow \mathrm{y}$ <br> else $x \leftarrow z$ | fsub fs, fb, fa <br> fsel $£ x, f s, f z, f y$ | $(3,4,5)$ |
| $\text { if } \mathrm{a}=\mathrm{b} \text { then } \mathrm{x} \leftarrow \mathrm{y}$ $\text { else } x \leftarrow z$ | ```fsub fs,fa,fb fsel fx,fs,fy,fz fneg fs,fs fsel fx,fs,fx,fz``` | $(4,5)$ |

1. The unoptimized program affects the VXSNAN bit of the FPSCR, and therefore may cause the system error handler to be invoked if the corresponding exception is enabled, while the optimized program does not affect this bit. This property of the optimized program is incompatible with the IEEE standard.
2. The optimized program gives the incorrect result if a is a NaN .
3. The optimized program gives the incorrect result if a and/or $b$ is a NaN (except that it may give the correct result in some cases for the minimum and maximum functions, depending on how those functions are defined to operate on NaNs ).
4. The optimized program gives the incorrect result if $a$ and $b$ are infinities of the same sign. (Here it is assumed that Invalid Operation Exceptions are disabled, in which case the result of the subtraction is a NaN . The analysis is more complicated if Invalid Operation Exceptions are enabled, because in that case the target register of the subtraction is unchanged.)
5. The optimized program affects the $O X, U X, X X$, and VXISI bits of the FPSCR, and therefore may cause the system error handler to be invoked if the corresponding exceptions are enabled, while the unoptimized program does not affect these bits. This property of the optimized program is incompatible with the IEEE standard.

### 4.6.10 Floating-Point Status and Control Register Instructions

Every Floating-Point Status and Control Register instruction synchronizes the effects of all floating-point instructions executed by a given processor. Executing a Floating-Point Status and Control Register instruction ensures that all floating-point instructions previously initiated by the given processor have completed before the Floating-Point Status and Control Register instruction is initiated, and that no subsequent floating-point instructions are initiated by the given processor until the Floating-Point Status and Control Register instruction has completed. In particular:
■ All exceptions that will be caused by the previously initiated instructions are recorded in the FPSCR before the Floating-Point Status and Control Register instruction is initiated.

- All invocations of the system floating-point enabled exception error handler that will be caused by the previously initiated instructions have occurred before the Floating-Point Status and Control Register instruction is initiated.

■ No subsequent floating-point instruction that depends on or alters the settings of any FPSCR bits is initiated until the Floating-Point Status and Control Register instruction has completed.
(Floating-point Storage Access instructions are not affected.)

The instruction descriptions in this section refer to "FPSCR fields," where FPSCR field $k$ is FPSCR bits 4xk:4xk+3.

## Move From FPSCR X-form



The contents of the FPSCR are placed into register FRT.
Special Registers Altered:
CR1
(if $\mathrm{Rc}=1$ )

## Move to Condition Register from FPSCR X-form

```
mcrfs BF,BFA
```

| 63 | BF | $/ /$ | BFA | $/ /$ |  | I/I |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 9 | 11 | 14 | 16 | 21 |

The contents of FPSCR $32: 63$ field BFA are copied to Condition Register field BF. All exception bits copied are set to 0 in the FPSCR. If the FX bit is copied, it is set to 0 in the FPSCR.

## Special Registers Altered:

CR field BF

| FX OX | (if BFA=0) |
| :--- | :--- |
| UX ZX XX VXSNAN | (if $\mathrm{BFA}=1$ ) |
| VXISI VXIDI VXZDZ VXIMZ | (if $\mathrm{BFA}=2$ ) |
| VXVC | (if $\mathrm{BFA}=3$ ) |
| VXSOFT VXSQRT VXCVI | (if BFA=5) |

## Move To FPSCR Field Immediate X-form

| mtfsfi <br> mtfsfi. | $\begin{aligned} & B F, U, W \\ & B F, U, W \end{aligned}$ |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $63$ | ${ }_{6} \mathrm{BF}$ | II I/I <br> 9 11 | W  <br> 15 U | ${ }^{1} 10$ | 134 | Rc 31 |

The value of the $U$ field is placed into FPSCR field $B F+8 \times(1-W)$.
$\mathrm{FPSCR}_{\mathrm{FX}}$ is altered only if $\mathrm{BF}=0$ and $\mathrm{W}=0$.

## Special Registers Altered:

$$
\text { FPSCR field } B F+8 \times(1-W)
$$

CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

mtfsfi serves as both a basic and an extended mnemonic. The Assembler will recognize a mtfsfi mnemonic with three operands as the basic form, and a mtfsfi mnemonic with two operands as the extended form. In the extended form the W operand is omitted and assumed to be 0 .

## Programming Note

When FPSCR $32: 35$ is specified, bits 32 (FX) and 35 (OX) are set to the values of $U_{0}$ and $U_{3}$ (i.e., even if this instruction causes $O X$ to change from 0 to 1 , $F X$ is set from $U_{0}$ and not by the usual rule that $F X$ is set to 1 when an exception bit changes from 0 to 1). Bits 33 and 34 (FEX and VX) are set according to the usual rule, given on page 125, and not from $\mathrm{U}_{1: 2}$.

## Move To FPSCR Fields XFL-form

| mtfsf mtfsf. | FLM,FRB,L,W FLM,FRB,L,W |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 063 | L  <br> 6 7 | FLM | $\begin{array}{\|l\|l} \hline W & \text { FRB } \\ 15 & 16 \end{array}$ | 21711 | Rc <br> 31 |

The FPSCR is modified as specified by the FLM, $L$, and W fields.

$$
L=0
$$

The contents of register FRB are placed into the FPSCR under control of the $W$ field and the field mask specified by FLM. W and the field mask identify the 4-bit fields affected. Let i be an integer in the range $0-7$. If $\mathrm{FLM}_{\mathrm{i}}=1$ then FPSCR field k is set to the contents of the corresponding field of register FRB, where $\mathrm{k}=\mathrm{i}+8 \times(1-\mathrm{W})$.

L = 1
The contents of register FRB are placed into the FPSCR.
$\mathrm{FPSCR}_{\text {FX }}$ is not altered implicitly by this instruction.

## Special Registers Altered:

FPSCR fields selected by mask, L, and W CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

mtfsf serves as both a basic and an extended mnemonic. The Assembler will recognize a mtfsf mnemonic with four operands as the basic form, and a mtfsf mnemonic with two operands as the extended form. In the extended form the W and L operands are omitted and both are assumed to be 0.

## Programming Note

Updating fewer than eight fields of the FPSCR may have substantially poorer performance on some implementations than updating eight fields or all of the fields.

## Programming Note

If $L=1$ or if $L=0$ and $F P S C R$ 32:35 is specified, bits 32 ( FX ) and 35 (OX) are set to the values of (FRB) 32 and (FRB) 35 (i.e., even if this instruction causes OX to change from 0 to $1, F X$ is set from $(F R B)_{32}$ and not by the usual rule that $F X$ is set to 1 when an exception bit changes from 0 to 1). Bits 33 and 34 ( $F E X$ and $V X$ ) are set according to the usual rule, given on page 125, and not from (FRB) 33:34. $^{\text {. }}$

Move To FPSCR Bit 0 X-form


Bit BT +32 of the $F P S C R$ is set to 0 .

## Special Registers Altered:

FPSCR bit BT+32
CR1 (if $R C=1)$

## Programming Note

Bits 33 and 34 (FEX and VX) cannot be explicitly reset.

## Move To FPSCR Bit 1 X-form

$\left.\begin{array}{l}\begin{array}{l}\text { mtfsb1 } \\ \text { mtfsb1. }\end{array} \text { BT } \\ \text { BT }\end{array} \quad \begin{array}{l}\text { (Rc=0) } \\ \text { (Rc=1) }\end{array}\right)$

Bit $B T+32$ of the $F P S C R$ is set to 1 .

## Special Registers Altered:

FPSCR bits BT +32 and $F X$
CR1
(if $R C=1$ )

## Programming Note

Bits 33 and 34 (FEX and VX) cannot be explicitly set.

# Chapter 5. Decimal Floating-Point 

### 5.1 Decimal Floating-Point (DFP) Facility Overview

This chapter describes the behavior of the decimal floating-point facility, the supported data types, formats, and classes, and the usage of registers. Also included are the execution model, exceptions, and instructions supported by the decimal floating-point facility.

The decimal floating-point (DFP) facility shares the 32 floating-point registers (FPRs) and the Floating-Point Status and Control Register (FPSCR) with the float-ing-point (BFP) facility. However, the interpretation of data formats in the FPRs, and the meaning of some control and status bits in the FPSCR are different between the BFP and DFP facilities.

The DFP facility also shares the Condition Register (CR) with the fixed-Point facility, the BFP faciltiy, and the vector facility.

The DFP facility supports three DFP data formats: DFP Short (single precision), DFP Long (double precision), and DFP Extended (quad precision). Most operations are performed on DFP Long or DFP Extended format directly. Support for DFP Short is limited to conversion to and from DFP Long. Some DFP instructions operate on other data types, including signed or unsigned binary fixed-point data, and signed or unsigned decimal data.

DFP instructions are provided to perform arithmetic, compare, test, quantum-adjustment, conversion, and format operations on operands held in FPRs or FPR pairs.

## - Arithmetic instructions

These instructions perform addition, subtraction, multiplication, and division operations.

- Compare instructions

These instructions perform a comparison operation on the numerical value of two DFP operands.

■ Test instructions

These instructions test the data class, the data group, the exponent, or the number of significant digits of a DFP operand.
■ Quantum-adjustment instructions
These instructions convert a DFP number to a result in the form that has the designated exponent, which may be explicitly or implicitly specified.

- Conversion instructions

These instructions perform conversion between different data formats or data types.

- Format instructions

These instructions facilitate composing or decomposing a DFP operand.

These instructions are described in Section 5.6 "DFP Instruction Descriptions" on page 194.

The three DFP data formats allow finite numbers to be represented with different precision and ranges. Special codes are also provided to represent +Infinity, -Infinity, Quiet NaN (Not-a-Number), and Signaling NaN . Operations involving infinities produce results obeying traditional mathematical conventions. NaNs have no mathematical interpretation. The encoding of NaNs provides a diagnostic information field. This diagnostic field may be used to indicate such things as the source of an uninitialized variable or the reason an invalid result was produced.

The DFP processor recognizes a set of DFP exceptions which are indicated via bits set in the FPSCR. Additionally, the DFP exception actions depend on the setting of the various exception enable bits in the FPSCR.

The following DFP exceptions are detected by the DFP processor. The exception status bits in the FPSCR are indicated in parentheses.

- Invalid Operation Exception
$\infty-\infty$
$\infty \div \infty$
$0 \div 0$
(VXZDZ)
$\infty \times 0$
(VXIMZ)
Invalid Compare
(VXVC)

```
        Invalid conversion
- Zero Divide Exception
- Overflow Exception
- Underflow Exception
- Inexact Exception
```

    (VXCVI)
    Each DFP exception and each category of Invalid Operation Exception has an exception status bit in the FPSCR. In addition, each of the five DFP exceptions has a corresponding enable bit in the FPSCR. These enable bits enable or disable the invocation of the system floating-point enabled exception error handler, and may affect the setting of some exception status bits in the FPSCR.
The usage of these bits by the DFP facility differs from the usage by the BFP facility. Section 5.5.10 "DFP Exceptions" on page 186 provides a detailed discussion of DFP exceptions, including the effects of the enable bits.

### 5.2 DFP Register Handling

The following sections describe first how the float-ing-point registers are utilized by the DFP facility. The subsequent section covers the DFP usage of CR and FPSCR.

### 5.2.1 DFP Usage of Floating-Point Registers

The DFP facility shares the same 32 64-bit FPRs with the BFP facility. Like the FP instructions, DFP instructions also use 5-bit fields for designating the FPRs to hold the source or target operands.

When data in DFP Short format is held in a FPR, it occupies the rightmost 32 bits of the FPR. The Load Floating-Point as Integer Word Algebraic instruction is provided to load the rightmost 32 bits of a FPR with a single-word data from storage. The Store Floating-Point as Integer Word instruction is available to store the rightmost 32 bits of a FPR to a storage location.
Data in DFP Long format, 64-bit binary fixed-point values, or 64-bit BCD values is held in a FPR using all 64 bits. Data of 64 bits may be loaded from storage via any of the Load Floating-Point Double instructions and stored via any of the Store Floating-Point Double instructions.

Data in DFP Extended format or 128-bit BCD values is held in an even-odd FPR pair using all 128 bits. Data of 128 bits must be loaded into the desired even-odd pair of floating-point registers using an appropriate sequence of the Load Floating-Point Double instructions and stored using an appropriate sequence of the Store Floating-Point Double instructions.
Data used as a source operand by any Decimal Float-ing-Point instruction that was produced, either directly
or indirectly, by a Load Floating-Point Single instruction, a Floating Round to Single-Precision instruction, or a binary floating-point single-precision arithmetic instruction is boundedly undefined.

When an even-odd FPR pair is used to hold a 128-bit operand, the even-numbered FPR is used to hold the leftmost doubleword of the operand and the next higher-numbered FPR is used to hold the rightmost doubleword. A DFP instruction designating an odd-numbered FPR for a 128-bit operand is an invalid instruction form.

## Programming Note

The Floating-Point Move instructions can be used to move operands between FPRs.

The bit definitions for the FPSCR are as follows.

## Bit(s) Description

0:28 Reserved
29:31 DFP Rounding Control (DRN)
See Section 5.5.2, "Rounding Mode Specification" on page 183.
000 Round to Nearest, Ties to Even
001 Round toward Zero
010 Round toward + Infinity
011 Round toward -Infinity
100 Round to Nearest, Ties away from 0
101 Round to Nearest, Ties toward 0
110 Round to away from Zero
111 Round to Prepare for Shorter Precision

## Programming Note

$\mathrm{FPSCR}_{28}$ is reserved for extension of the DRN field, therefore DRN may be set using the mtfsfi instruction to set the rounding mode.

Floating-Point Exception Summary (FX)
Every floating-point instruction, except mtfsfi and $\boldsymbol{m t f s f}$, implicitly sets FPSCR $_{F X}$ to 1 if that instruction causes any of the floating-point exception bits in the FPSCR to change from 0 to 1. merfs, mtfsfi, mtfsf, mtfsbO, and mtfsb1 can alter FPSCR $_{\text {FX }}$ explicitly.
33 Floating-Point Enabled Exception Summary (FEX)
This bit is the OR of all the floating-point exception bits masked by their respective enable bits. mcrfs, mtfsfi, mtfsf, mtfsbO, and mtfsb1 cannot alter FPSCR FEX explicitly.
34 Floating-Point Invalid Operation Exception Summary (VX)
This bit is the OR of all the Invalid Operation exception bits. mcrfs, mtfsfi, mtfsf, mtfsbO, and $\boldsymbol{m t f s b 1}$ cannot alter FPSCR $_{\mathrm{Vx}}$ explicitly.

Floating-Point Overflow Exception (OX)See Section 5.5.10.3, "Overflow Exception" on page 189.
$41 \quad$ Floating-Point Invalid Operation Exception
( $\infty \div \infty$ ) (VXIDI)
See Section 5.5.10.1.
Floating-Point Invalid Operation Exception ( $0 \div 0$ ) (VXZDZ)
See Section 5.5.10.1.
Floating-Point Invalid Operation Exception ( $\infty \times 0$ ) (VXIMZ)
See Section 5.5.10.1.
Floating-Point Invalid Operation Exception (Invalid Compare) (VXVC)
See Section 5.5.10.1.
Floating-Point Fraction Rounded (FR)
The last Arithmetic or Rounding and Conversion instruction incremented the fraction during rounding. See Section 5.5.1, "Rounding" on page 182. This bit is not sticky.
Floating-Point Fraction Inexact (FI)
The last Arithmetic or Rounding and Conversion instruction either produced an inexact result during rounding or caused a disabled Overflow Exception. See Section 5.5.1. This bit is not sticky.

- If the instruction affects FPSCR $_{\text {FI }}$, the new value of FPSCR $_{X X}$ is obtained by ORing the old value of FPSCR $_{X X}$ with the new value of FPSCR ${ }_{\mathrm{FI}}$.
- If the instruction does not affect FPSCR $_{\text {Fl }}$, the value of FPSCR $_{X X}$ is unchanged.
Floating-Point Invalid Operation Exception (SNaN) (VXSNAN)
See Section 5.5.10.1, "Invalid Operation Exception" on page 188.

Floating-Point Invalid Operation Exception ( $\infty$ - $\infty$ ) (VXISI) See Section 5.5.10.1.

See the definition of FPSCR $_{X X}$, above, regarding the relationship between FPSCR $_{\text {FI }}$ and FPSCR $_{X X}$.
Floating-Point Result Flags (FPRF)
This field is set as described below. For arithmetic, rounding, and conversion instructions, the field is set based on the result placed into the target register, except that if any portion of the result is undefined then the value placed into FPRF is undefined.

Floating-Point Result Class Descriptor (C)
Arithmetic, rounding, and conversion instructions may set this bit with the FPCC bits, to indicate the class of the result as shown in Figure 61 on page 178.

Floating-Point Condition Code (FPCC)
Floating-point Compare and DFP Test instructions set one of the FPCC bits to 1 and the other three FPCC bits to 0 . Arithmetic, rounding, and conversion instructions may set the FPCC bits with the C bit, to indicate the class of the result as shown in Figure 61 on page 178. Note that in this case the high-order three bits of the FPCC retain their relational significance indicating that the value is less than, greater than, or equal to zero.
Floating-Point Less Than or Negative (FL or <)

Floating-Point Greater Than or Positive (FG or >)

Floating-Point Equal or Zero (FE or =)
Floating-Point Unordered or NaN (FU or ?)
Reserved
Floating-Point Invalid Operation Exception (Software Request) (VXSOFT)
This bit can be altered only by mcrfs, mtfsfi, mtfsf, mtfsb0, or mtfsb1. See Section 5.5.10.1, "Invalid Operation Exception" on page 188.
Neither used nor changed by DFP.

## Programming Note

Although the architecture does not provide a DFP square root instruction, if software simulates such an instruction, it should set bit 54 whenever the source operand of the square root function is invalid.

Floating-Point Invalid Operation Exception (Invalid Conversion) (VXCVI)
See Section 5.5.10.1.
Floating-Point Invalid Operation Exception Enable (VE)
See Section 5.5.10.1.

57 Floating-Point Overflow Exception Enable (OE)
See Section 5.5.10.3, "Overflow Exception" on page 189.

58 Floating-Point Underflow Exception Enable (UE)
See Section 5.5.10.4, "Underflow Exception" on page 190.
59 Floating-Point Zero Divide Exception Enable (ZE)
See Section 5.5.10.2, "Zero Divide Exception" on page 189.
60 Floating-Point Inexact Exception Enable (XE)
See Section 5.5.10.5, "Inexact Exception" on page 191

61 Reserved (not used by DFP)
62:63 Binary Floating-Point Rounding Control (RN)
See Section 5.5.1, "Rounding" on page 182.
00 Round to Nearest
01 Round toward Zero
10 Round toward + Infinity
11 Round toward -Infinity

| Result Flags | Result Value Class |
| :---: | :---: |
| C < > = ? |  |
| 00001 | Signaling NaN (DFP only) |
| 10001 | Quiet NaN |
| 01001 | - Infinity |
| 01000 | - Normal Number |
| 11000 | - Subnormal Number |
| 10010 | - Zero |
| 00010 | + Zero |
| 10100 | + Subnormal Number |
| 00100 | + Normal Number |
| 00101 | + Infinity |

Figure 61. Floating-Point Result Flags

### 5.3 DFP Support for Non-DFP Data Types

In addition to the DFP data types, the DFP processor provides limited support for the following non-DFP data types: signed or unsigned binary fixed-point data, and signed or unsigned decimal data.
In unsigned binary fixed-point data, all bits are used to express the absolute value of the number. For signed binary fixed-point data, the leftmost bit represents the sign, which is followed by the numeric field. Positive numbers are represented in true binary notation with the sign bit set to zero. When the value is zero, all bits
are zeros, including the sign bit. Negative numbers are represented in two's complement binary notation with a one in the sign-bit position.
For decimal data, each byte contains a pair of four-bit nibbles; each four-bit nibble contains a binary-coded-decimal (BCD) code. There are two kinds of BCD codes: digit code and sign code. For unsigned decimal data, all nibbles contain a digit code (D) as shown in Figure 62

| D | D | D | D | $\ldots$ | D | D | D | D |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

Figure 62. Format for Unsigned Decimal Data
For signed decimal data, the rightmost nibble contains a sign code ( S ) and all other nibbles contain a digit code as shown in Figure 63.

| D | D | D | D | $\ldots$ | D | D | D | S |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

Figure 63. Format for Signed Decimal Data
The decimal digits 0-9 have the binary encoding 0000-1001. The preferred plus-sign codes are 1100 and 1111. The preferred minus sign code is 1101. These are the sign codes generated for the results of the Decode DPD To BCD instruction. A selection is provided by this instruction to specify which of the two preferred plus sign codes is to be generated. Alternate sign codes are also recognized as valid in the sign position: 1010 and 1110 are alternate sign codes for plus, and 1011 is an alternate sign code for minus. Alternate sign codes are accepted for any source operand, but are not generated as a result by the instruction. When an invalid digit or sign code is detected by the Encode $B C D$ To DPD instruction, an invalid-opera-
tion exception occurs. A summary of digit and sign codes are provided in Figure 64.

| Binary <br> Code | Recognized As |  |
| :---: | :---: | :---: |
|  | Digit | Sign |
| 0000 | 0 | Invalid |
| 0001 | 1 | Invalid |
| 0010 | 2 | Invalid |
| 0011 | 3 | Invalid |
| 0100 | 4 | Invalid |
| 0101 | 5 | Invalid |
| 0110 | 6 | Invalid |
| 0111 | 7 | Invalid |
| 1000 | 8 | Invalid |
| 1001 | 9 | Invalid |
| 1010 | Invalid | Plus |
| 1011 | Invalid | Minus |
| 1100 | Invalid | Plus (preferred; option 1) |
| 1101 | Invalid | Minus (preferred) |
| 1110 | Invalid | Plus |
| 1111 | Invalid | Plus (preferred; option 2) |

Figure 64. Summary of BCD Digit and Sign Codes

### 5.4 DFP Number Representation

A DFP finite number consists of three components: a sign bit, a signed exponent, and a significand. The signed exponent is a signed binary integer. The significand consists of a number of decimal digits, which are to the left of the implied decimal point. The rightmost digit of the significand is called the units digit. The numerical value of a DFP finite number is represented as $(-1)^{\text {sign }} \times$ significand $\times 10^{\text {exponent }}$ and the unit value of this number is ( $\left.1 \times 10^{\text {exponent }}\right)$, which is called the quantum.
DFP finite numbers are not normalized. This allows leading zeros and trailing zeros to exist in the significand. This unnormalized DFP number representation allows some values to have redundant forms; each form represents the DFP number with a different combination of the significand value and the exponent value. For example, $1000000 \times 10^{5}$ and $10 \times 10^{10}$ are two different forms of the same numerical value. A form of this number representation carries information about both the numerical value and the quantum of a DFP finite number.

The significant digits of a DFP finite number are the digits in the significand beginning with the leftmost nonzero digit and ending with the units digit.

### 5.4.1 DFP Data Format

DFP numbers and NaNs may be represented in FPRs in any of the three data formats: DFP Short, DFP Long, or DFP Extended. The contents of each data format represent encoded information. Special codes are assigned to NaNs and infinities. Different formats support different sizes in both significand and exponent. Arithmetic, compare, test, quantum-adjustment, and format instructions are provided for DFP Long and DFP Extended formats only.

The sign is encoded as a one bit binary value. Significand is encoded as an unsigned decimal integer in two distinct parts. The leftmost digit (LMD) of the significand is encoded as part of the combination field; the remaining digits of the significand are encoded in the trailing significand field. The exponent is contained in the combination field in two parts. However, prior to encoding, the exponent is converted to an unsigned binary value called the biased exponent by adding a bias value which is a constant for each format. The two leftmost bits of the biased exponent are encoded with the leftmost digit of the significand in the leftmost bits of the combination field. The rest of the biased exponent occupies the remaining portion of the combination field.

### 5.4.1.1 Fields Within the Data Format

The DFP data representation comprises three fields, as diagrammed below for each of the three formats:


Figure 65. DFP Short format

| $S$ | $G$ |  | $T$ |
| :--- | :--- | :--- | :--- |
| 0 | 1 | 14 | 63 |

Figure 66. DFP Long format

| $S$ | $G$ |  | $T$ |
| :--- | :--- | :--- | ---: |
| 01 | 18 | 63 |  |
|  | T (continued) |  |  |
| 64 |  | 127 |  |

Figure 67. DFP Extended format
The fields are defined as follows:
Sign bit (S)
The sign bit is in bit 0 of each format, and is zero for plus and one for minus.
Combination field (G)
As the name implies, this field provides a combination of the exponent and the left-most digit (LMD) of the significand, for finite numbers, or provides a special code
for denoting the value as either a Not-a-Number or an Infinity.
The first 5 bits of the combination field contain the encoding of NaN or infinity, or the two leftmost bits of the biased exponent and the leftmost digit (LMD) of the significand. The following tables show the encoding:

| $\mathbf{G}_{\mathbf{0}: \mathbf{4}}$ | Description |
| :---: | :--- |
| 11111 | NaN |
| 11110 | Infinity |
| All others | Finite Number (see Figure 69) |

Figure 68. Encoding of the $G$ field for Special Symbols

| LMD | Leftmost 2-bits of biased exponent |  |  |
| :---: | :---: | :---: | :---: |
|  | $\mathbf{0 0}$ | $\mathbf{0 1}$ | $\mathbf{1 0}$ |
| 0 | 00000 | 01000 | 10000 |
| 1 | 00001 | 01001 | 10001 |
| 2 | 00010 | 01010 | 10010 |
| 3 | 00011 | 01011 | 10011 |
| 4 | 00100 | 01100 | 10100 |
| 5 | 00101 | 01101 | 10101 |
| 6 | 00110 | 01110 | 10110 |
| 7 | 00111 | 01111 | 10111 |
| 8 | 11000 | 11010 | 11100 |
| 9 | 11001 | 11011 | 11101 |

Figure 69. Encoding of bits $0: 4$ of the $\mathbf{G}$ field for Finite Numbers

For DFP finite numbers, the rightmost N-5 bits of the N -bit combination field contain the remaining bits of the biased exponent. For NaNs, bit 5 of the combination field is used to distinguish a Quiet NaN from a Signaling NaN ; the remaining bits in a source operand are ignored and they are set to zeros in a target operand by most operations. For infinities, the rightmost N-5 bits of the N -bit combination field of a source operand are ignored and they are set to zeros in a target operand by most operations.

## Trailing Significand field (T)

For DFP finite numbers, this field contains the remaining significand digits. For NaNs, this field may be used to contain diagnostic information. For infinities, contents in this field of a source operand are ignored and they are set to zeros in a target operand by most operations. The trailing significand field is a multiple of 10-bit blocks. The multiple depends on the format. Each 10-bit block is called a declet and represents three decimal digits, using the Densely Packed Decimal (DPD) encoding defined in Appendix B.

### 5.4.1.2 Summary of DFP Data Formats

The properties of the three DFP formats are summarized in the following table:.

|  | Format |  |  |
| :---: | :---: | :---: | :---: |
|  | DFP Short | DFP Long | DFP Extended |
| Widths (bits): |  |  |  |
| Format | 32 | 64 | 128 |
| Sign (S) | 1 | 1 | 1 |
| Combination (G) | 11 | 13 | 17 |
| Trailing Significand (T) | 20 | 50 | 110 |
| Exponent: |  |  |  |
| Maximum biased | 191 | 767 | 12,287 |
| Maximum ( $\mathrm{X}_{\text {max }}$ ) | 90 | 369 | 6111 |
| Minimum ( $\mathrm{X}_{\text {min }}$ ) | -101 | -398 | -6176 |
| Bias | 101 | 398 | 6176 |
| Precision (p) (digits) | 7 | 16 | 34 |
| Magnitude: |  |  |  |
| Maximum normal number ( $\mathrm{N}_{\mathrm{max}}$ ) | $\left(10^{7}-1\right) \times 10^{90}$ | $\left(10^{16}-1\right) \times 10^{369}$ | $\left(10^{34}-1\right) \times 10^{6111}$ |
| Minimum normal number ( $\mathrm{N}_{\text {min }}$ ) | $1 \times 10^{-95}$ | $1 \times 10^{-383}$ | $1 \times 10^{-6143}$ |


|  | Format |  |  |
| :---: | :---: | :---: | :---: |
|  | DFP Short | DFP Long | DFP Extended |
| Minimum subnormal number $\left(\mathrm{D}_{\min }\right)$ | $1 \times 10^{-101}$ | $1 \times 10^{-398}$ | $1 \times 10^{-6176}$ |

Figure 70. Summary of DFP Formats

### 5.4.1.3 Preferred DPD Encoding

Execution of DFP instructions decodes source operands from DFP data formats to an internal format for processing, and encodes the operation result before the final result is returned as the target operand.

As part of the decoding process, declets in the trailing significand field of source operands are decoded to their corresponding $B C D$ digit codes using the DPD-to-BCD decoding algorithm. As part of the encoding process, BCD digit codes to be stored into the trailing significand field of the target operand are encoded into declets using the BCD-to-DPD encoding algorithm. Both the decoding and encoding algorithms are defined in Appendix B.
As explained in Appendix B, there are eight 3-digit decimal values that have redundant DPD codes and one preferred DPD code. All redundant DPD codes are recognized in source operands for the associated 3-digit decimal number. DFP operations will always generate the preferred DPD codes for the trailing significand field of the target operand.

### 5.4.2 Classes of DFP Data

There are six classes of DFP data, which include numerical and nonnumeric entities. The numerical entities include zero, subnormal number, normal number, and infinity data classes. The nonnumeric entities include quiet and signaling NaNs data classes. The value of a DFP finite number, including zero, subnormal number, and normal number, is a quantization of the real number based on the data format. The Test Data Class instruction may be used to determine the class of a DFP operand. In general, an operation that returns a DFP result sets the FPSCR FPRF field to indicate the data class of the result.

The following tables show the value ranges for finite-number data classes, and the codes for NaNs and infinities.

| Data Class | Sign | Magnitude |
| :--- | :---: | :---: |
| Zero | $\pm$ | $0^{\star}$ |
| Subnormal | $\pm$ | $\mathrm{D}_{\min } \leq \mid \mathrm{XI}<\mathrm{N}_{\min }$ |
| Normal | $\pm$ | $\mathrm{N}_{\min } \leq\|\mathrm{Y}\| \leq \mathrm{N}_{\max }$ |

* The significand is zero and the exponent is any representable value

Figure 71. Value Ranges for Finite Number Data Classes

| Data Class | S | G | T |
| :--- | :---: | :---: | :---: |
| + Infinity | 0 | $11110 x x x \ldots x x x$ | $x x x \ldots x x x$ |
| - Infinity | 1 | $11110 x x x \ldots x x x$ | $x x x \ldots x x x$ |
| Quiet NaN | $x$ | $111110 x x \ldots x x x$ | $x x x \ldots x x x$ |
| Signaling NaN | $x$ | $111111 x x \ldots x x x$ | $x x x \ldots x x x$ |
| x Don't care |  |  |  |
|  |  |  |  |

Figure 72. Encoding of NaN and Infinity Data Classes

## Zeros

Zeros have a zero significand and any representable value in the exponent. A +0 is distinct from -0 , and zeros with different exponents are distinct, except that comparison treats them as equal.

## Subnormal Numbers

Subnormal numbers have values that are smaller than $\mathrm{N}_{\text {min }}$ and greater than zero in magnitude.

## Normal Numbers

Normal numbers are nonzero finite numbers whose magnitude is between $N_{\text {min }}$ and $N_{\text {max }}$ inclusively.

## Infinities

Infinities are represented by $0 b 11110$ in the leftmost 5 bits of the combination field. When an operation is defined to generate an infinity as the result, a default infinity is sometimes supplied. A default infinity has all remaining bits in the combination field and trailing significand field set to zeros.

When infinities are used as source operands, only the leftmost 5 bits of the combination field are interpreted (i.e., Ob11110 indicates the value is an infinity). The trailing significand field of infinities is usually ignored. For generated infinities, the leftmost 5 bits of the combination field are set to 0 b11110 and all remaining combination bits are set to zero.

Infinities can participate in most arithmetic operations and give a consistent result. In comparisons, any +Infinity compares greater than any finite number, and any -Infinity compares less than any finite number. All +Infinity are compared equal and all -Infinity are compared equal.

## Signaling and Quiet NaNs

There are two types of Not-a-Numbers (NaNs), Signaling (SNaN) and Quiet (QNaN).

Ob111110 in the leftmost 6 bits of the combination field indicates a Quiet NaN, whereas Ob111111 indicates a Signaling NaN .
A special QNaN is sometimes supplied as the default $Q N a N$ for a disabled invalid-operation exception; it has a plus sign, the leftmost 6 bits of the combination field set to $0 b 111110$ and remaining bits in the combination field and the trailing significand field set to zero.

Normally, source QNaNs are propagated during operations so that they will remain visible at the end. When a QNaN is propagated, the sign is preserved, the decimal value of the trailing significand field is preserved but reencoded using the preferred DPD codes, and the contents in the rightmost $\mathrm{N}-6$ bits of the combination field set to zero, where N is the width of the combination field for the format.

A source SNaN generally causes an invalid-operation exception. If the exception is disabled, the SNaN is converted to the corresponding QNaN and propagated. The primary encoding difference between an SNaN and a QNaN is that bit 5 of an SNaN is 1 and bit 5 of a QNaN is 0 . When an SNaN is propagated as a QNaN, bit 5 is set to 0 , and, just as with QNaN proagation, the sign is preserved, the decimal value of the trailing significand field is preserved but reencoded using the preferred DPD codes, and the contents in the rightmost $\mathrm{N}-6$ bits of the combination field set to zero, where N is the width of the combination field for the format. For some format-conversion instructions, a source SNaN does not cause an invalid-operation exception, and an SNaN is returned as the target operand.

For instructions with two source NaNs and a NaN is to be propagated as the result, do the following.

- If there is a QNaN in FRA and an SNaN in FRB, the SNaN in FRB is propagated.
- Otherwise, propagate the NaN is FRA.


### 5.5 DFP Execution Model

DFP operations are performed as if they first produce an intermediate result correct to infinite precision and with unbounded range. The intermediate result is then rounded to the destination's precision according to one of the eight DFP rounding modes. If the rounded result has only one form, it is delivered as the final result; if the rounded result has redundant forms, then an ideal exponent is used to select the form of the final result. The ideal exponent determines the form, not the value, of the final result. (See Section 5.5.3 "Formation of Final Result" on page 184.)

### 5.5.1 Rounding

Rounding takes a number regarded as infinitely precise and, if necessary, modifies it to fit the destination's precision. The destination's precision of an operation defines the set of permissible resultant values. For
most operations, the destination's precision is the tar-get-format precision and the permissible resultant values are those values representable in the target format. For some special operations, the destination precision is constrained by both the target format and some additional restrictions, and the permissible resultant values are a subset of the values representable in the target format.

Rounding sets FPSCR bits FR and FI. When an inexact exception occurs, Fl is set to one; otherwise, Fl is set to zero. When an inexact exception occurs and if the rounded result is greater in magnitude than the intermediate result, then FR is set to one; otherwise, FR is set to zero. The exception is the Round to FP Integer Without Inexact instruction, which always sets FR and FI to zero. Rounding may cause an overflow exception or underflow exception; it may also cause an inexact exception.

Refer to Figure 73 below for rounding. Let $Z$ be the intermediate result of a DFP operation. Z may or may not fit in the destination's precision. If $Z$ is exactly one of the permissible representable resultant values, then the final result in all rounding modes is $Z$. Otherwise, either Z1 or Z2 is chosen to approximate the result, where Z1 and Z 2 are the next larger and smaller permissible resultant values, respectively.


Figure 73. Rounding

## Round to Nearest, Ties to Even

Choose the value that is closer to Z ( Z 1 or Z 2 ). In case of a tie, choose the one whose units digit would have been even in the form with the largest common quantum of the two permissible resultant values. However, an infinitely precise result with magnitude at least ( $\mathrm{N}_{\max }$ $+0.5 \mathrm{Q}\left(\mathrm{N}_{\text {max }}\right)$ ) is rounded to infinity with no change in sign; where $\mathrm{Q}\left(\mathrm{N}_{\max }\right)$ is the quantum of $\mathrm{N}_{\text {max }}$.

## Round toward 0

Choose the smaller in magnitude (Z1 or Z2).
Round toward $+\infty$
Choose Z1.
Round toward $-\infty$
Choose Z2.
Round to Nearest, Ties away from 0
Choose the value that is closer to Z ( Z 1 or Z 2 ). In case
of a tie, choose the larger in magnitude (Z1 or Z2). However, an infinitely precise result with magnitude at least $\left(N_{\max }+0.5 Q\left(N_{\max }\right)\right)$ is rounded to infinity with no change in sign; where $\mathrm{Q}\left(\mathrm{N}_{\max }\right)$ is the quantum of $\mathrm{N}_{\max }$.

## Round to Nearest, Ties toward 0

Choose the value that is closer to $Z(Z 1$ or $Z 2)$. In case of a tie, choose the smaller in magnitude (Z1 or Z2). However, an infinitely precise result with magnitude greater than $\left(\mathrm{N}_{\max }+0.5 \mathrm{Q}\left(\mathrm{N}_{\max }\right)\right)$ is rounded to infinity with no change in sign; where $Q\left(N_{\max }\right)$ is the quantum of $\mathrm{N}_{\text {max }}$.

## Round away from 0

Choose the larger in magnitude (Z1 or Z2).
Round to prepare for shorter precision
Choose the smaller in magnitude (Z1 or Z2). If the selected value is inexact and the units digit of the selected value is either 0 or 5 , then the digit is incremented by one and the incremented result is delivered. In all other cases, the selected value is delivered. When a value has redundant forms, the units digit is determined by using the form that has the smallest exponent.

### 5.5.2 Rounding Mode Specification

Unless otherwise specified in the instruction definition, the rounding mode used by an operation is specified in the DFP rounding control (DRN) field of the FPSCR. The eight DFP rounding modes are encoded in the DRN field as specified in the table below.

```
DRN Rounding Mode
000 Round to Nearest, Ties to Even
0 0 1 ~ R o u n d ~ t o w a r d ~ 0
0 1 0 ~ R o u n d ~ t o w a r d ~ + I n f i n i t y ~
011 Round toward -Infinity
100 Round to Nearest, Ties away from 0
101 Round to Nearest, Ties toward 0
110 Round away from 0
111 Round to Prepare for Shorter Precision
```

Figure 74. Encoding of DFP Rounding-Mode Control (DRN)

For the quantum-adjustment, a 2-bit immediate field, called RMC (Rounding Mode Contro), in the instruction specifies the rounding mode used. The RMC field may contain a primary encoding or a secondary encoding. For Quantize, Quantize Immediate, and Reround, the RMC field contains the primary encoding. For Round to FP Integer the field contains either encoding, depending on the setting of a RMC-encoding-selection
bit. The following tables define the primary encoding and the secondary encoding.

| Primary | Rounding Mode |
| :---: | :--- |
| RMC | Round to nearest, ties to even |
| 00 | Round toward 0 |
| 01 | Round to nearest, ties away from 0 |
| 10 | Round according to FPSCR |
| 11 | DRN |

Figure 75. Primary Encoding of Rounding-Mode Control

## Secondary RMC <br> Rounding Mode <br> 00 <br> Round to $+\infty$ <br> 01 <br> Round to - $\infty$ <br> 10 Round away from 0 <br> 11 Round to nearest, ties toward 0

Figure 76. Secondary Encoding of Rounding-Mode Control

### 5.5.3 Formation of Final Result

An ideal exponent is defined for each DFP instruction that returns a DFP data operand.

### 5.5.3.1 Use of Ideal Exponent

For all DFP operations,
■ if the rounded intermediate result has only one form, then that form is delivered as the final result.

- if the rounded intermediate result has redundant. forms and is exact, then the form with the exponent closest to the ideal exponent is delivered.
- if the rounded intermediate result has redundant forms and is inexact, then the form with the smallest exponent is delivered.

The following table specifies the ideal exponent for each instruction.

| Operations | Ideal Exponent |
| :---: | :---: |
| Add | $\min (\mathrm{E}(\mathrm{FRA}), \mathrm{E}(\mathrm{FRB})$ ) |
| Subtract | $\min (\mathrm{E}(\mathrm{FRA}), \mathrm{E}(\mathrm{FRB}))$ |
| Multiply | $E(F R A)+E(F R B)$ |
| Divide | E(FRA) - E(FRB) |
| Quantize-Immediate | See Instruction Description |
| Quantize | E(FRA) |
| Reround | See Instruction Description |
| Round to FP Integer | $\max (0, \mathrm{E}$ (FRA) ) |
| Convert to DFP Long | E(FRA) |
| Convert to DFP Extended | E(FRA) |
| Round to DFP Short | E(FRA) |
| Round to DFP Long | E(FRA) |
| Convert from Fixed | 0 |
| Encode BCD to DPD | 0 |
| Insert Biased Exponent | E(FRA) |
| Notes: <br> $E(x)$ - exponent of | e DFP operand in register $x$ |

Figure 77. Summary of Ideal Exponents

### 5.5.4 Arithmetic Operations

Four arithmetic operations are provided: Add, Subtract, Multiply, and Divide.

### 5.5.4.1 Sign of Arithmetic Result

The following rules govern the sign of an arithmetic operation when the operation does not yield an exception. They apply even when the operands or results are zeros or infinities.
■ The sign of the result of an add operation is the sign of the source operand having the larger absolute value. If both source operands have the same sign, the sign of the result of an add operation is the same as the sign of the source operands. When the sum of two operands with opposite signs is exactly zero, the sign of the result is positive in all rounding modes except Round toward $-\infty$, in which case the sign is negative.

- The sign of the result of the subtract operation $x-y$ is the same as the sign of the result of the add operation $x+(-y)$.
- The sign of the result of a multiply or divide operation is the exclusive-OR of the signs of the source operands.


### 5.5.5 Compare Operations

Two sets of instructions are provided for comparing numerical values: Compare Ordered and Compare Unordered. In the absence of NaNs , these instructions work the same. These instructions work differently when either of the followings is true:

1. At least one source operand of the instruction is an SNaN and the invalid-operation exception is disabled.
2. When there is no SNaN in any source operand, at least one source operand of the instruction is a QNaN
In case 1, Compare Unordered recognizes an invalid-operation exception and sets the FPSCR ${ }_{\text {VXSNAN }}$ flag, but Compare Ordered recognizes the exception and sets both the FPSCR ${ }_{\text {VXSNAN }}$ and FPSCR $_{\text {Vxvc }}$ flags. In case 2, Compare Unordered does not recognize an exception, but Compare Ordered recognizes an invalid-operation exception and sets the FPSCR $_{\text {Vxvc }}$ flag.
For finite numbers, comparisons are performed on values, that is, all redundant forms of a DFP number are treated equal.
Comparisons are always exact and cannot cause an inexact exception.

Comparison ignores the sign of zero, that is, +0 equals -0.

Infinities with like sign compare equal, that is, $+\infty$ equals $+\infty$, and $-\infty$ equals $-\infty$.

A NaN compares as unordered with any other operand, whether a finite number, an infinity, or another NaN , including itself.

Execution of a compare instruction always completes, regardless of whether any DFP exception occurs or not, and whether the exception is enabled or not.

### 5.5.6 Test Operations

Four kinds of test operations are provided: Test Data Class, Test Data Group, Test Exponent, and Test Significance.

The Test Data Class instruction examines the contents of a source operand and determines if the operand is one of the specified data classes. The test result and the sign of the source operand are indicated in the FPSCR $_{\text {FPCC }}$ field and CR field BF.

The Test Data Group instruction examines the contents of a source operand and determines if the operand is one of the specified data groups. The test result and the sign of the source operand are indicated in the FPSCR $_{\text {FPCC }}$ field and CR field BF.
The Test Exponent instruction compares the exponent of the two source operands. The test operation ignores
the sign and significand of operands. Infinities compare equal, and NaNs compare equal. The test result is indicated in the FPSCR FPCC field and CR field BF.
The Test Significance instruction compares the number of significant digits of one source operand with the referenced number of significant digits in another source operand. The test result is indicated in the FPSCR FPCC field and CR field BF.

Execution of a test instruction does not cause any DFP exception.

### 5.5.7 Quantum Adjustment Operations

Four kinds of quantum-adjustment operations are provided: Quantize, Quantize Immediate, Reround, and Round To FP Integer. Each of them has an immediate field which specifies whether the rounding mode in FPSCR or a different one is to be used.
The Quantize instruction is used to adjust a DFP number to the form that has the specified target exponent. The Quantize Immediate instruction is similar to the Quantize instruction, except that the target exponent is specified in a 5 -bit immediate field as a signed binary integer and has a limited range.
The Reround instruction is used to simulate a DFP operation of a precision other than that of DFP Long or DFP Extended. For the Reround instruction to produce a result which accurately reflects that which would have resulted from a DFP operation of the desired precision $d$ in the range $\{1: 33\}$ inclusively, the following conditions must be met:
■ The precision of the preceding DFP operation must be at least one digit larger than $d$.

- The rounding mode used by the preceding DFP operation must be round-to-pre-pare-for-shorter-precision.

The Round To FP Integer instruction is used to round a DFP number to an integer value of the same format. The target exponent is implicitly specified, and is greater than or equal to zero.

### 5.5.8 Conversion Operations

There are two kinds of conversion operations: data-format conversion and data-type conversion.

### 5.5.8.1 Data-Format Conversion

The instructions Convert To DFP Long and Convert To DFP Extended convert DFP operands to wider formats; the instructions Round To DFP Short and Round To DFP Long convert DFP operands to narrower formats.
When converting a finite number to a wider format, the result is exact. When converting a finite number to a
narrower format, the source operand is rounded to the target-format precision, which is specified by the instruction, not by the target register size.
When converting a finite number, the ideal exponent of the result is the source exponent.

Conversion of an infinity or NaN to a different format does not preserve the source combination field. Let N be the width of the target format's combination field.

- When the result is an infinity or a QNaN, the contents of the rightmost $\mathrm{N}-5$ bits of the N -bit target combination field are set to zero.
- When the result is an SNaN , bit 5 of the target format's combination field is set to one and the rightmost $\mathrm{N}-6$ bits of the N -bit target combination field are set to zero.
When converting a NaN to a wider format or when converting an infinity from DFP Short to DFP Long, digits in the source trailing significand field are reencoded using the preferred DPD codes with sufficient zeros appended on the left to form the target trailing significand field. When converting a NaN to a narrower format or when converting an infinity from DFP Long to DFP Short, the appropriate number of leftmost digits of the source trailing significand field are removed and the remaining digits of the field are reencoded using the preferred DPD codes to form the target trailing significand field.

When converting an infinity between DFP Long and DFP Extended, a default infinity with the same sign is produced.
When converting an SNaN between DFP Short and DFP Long, it is converted to an SNaN without causing an invalid-operation exception. When converting an SNaN between DFP Long and DFP Extended, the invalid-operation exception occurs; if the invalid-operation exception is disabled, the result is converted to the corresponding QNaN.

### 5.5.8.2 Data-Type Conversion

The instructions Convert From Fixed and Convert To Fixed are provided to convert a number between the DFP data type and the signed 64-bit binary-integer data type.

Conversion of a signed 64-bit binary integer to a DFP Extended number is always exact.

Conversion of a DFP number to a signed 64-bit binary integer results in an invalid-operation exception when the converted value does not fit into the target format, or when the source operand is an infinity or NaN . When the exception is disabled, the most positive integer is returned if the source operand is a positive number or $+\infty$, and the most negative integer is returned if the source operand is a negative number, $-\infty$, or NaN .

### 5.5.9 Format Operations

The format instructions are provided to facilitate composing or decomposing a DFP number, and consist of Encode BCD To DPD, Decode DPD To BCD, Extract Biased Exponent, Insert Biased Exponent, Shift Significand Left Immediate, and Shift Significand Right Immediate. A source operand of SNaN does not cause an invalid-operation exception, and an SNaN may be produced as the target operand.

### 5.5.10 DFP Exceptions

This architecture defines the following DFP exceptions:

$$
\begin{aligned}
& \text { Invalid Operation Exception } \\
& \text { SNaN } \\
& \infty-\infty \\
& \infty \div \infty \\
& 0 \div 0 \\
& \infty \times 0 \\
& \text { Invalid Compare } \\
& \text { Invalid Conversion } \\
& \text { ■ Zero Divide Exception } \\
& \text { ■ Overflow Exception } \\
& \text { ■ Underflow Exception } \\
& \text { ■ Inexact Exception }
\end{aligned}
$$

These exceptions may occur during execution of a DFP instruction.

Each DFP exception, and each category of the Invalid Operation Exception, has an exception status bit in the FPSCR. In addition, each DFP exception has a corresponding enable bit in the FPSCR. The exception status bit indicates occurrence of the corresponding exception. If an exception occurs, the corresponding enable bit governs the result produced by the instruction and, in conjunction with the FE0 and FE1 bits (see the discussion of FE0 and FE1 below), whether and how the system floating-point enabled exception error handler is invoked. (In general, the enabling specified by the enable bit is of invoking the system error handler, not of permitting the exception to occur. The occurrence of an exception depends only on the instruction and its source operands, not on the setting of any control bits. The only deviation from this general rule is that the occurrence of an Underflow Exception may depend on the setting of the enable bit.)

A single instruction, other than mtfsfi or mtfsf, may set more than one exception bit only in the following cases:

- Inexact Exception may be set with Overflow Exception.
- Inexact Exception may be set with Underflow Exception.
- Invalid Operation Exception (SNaN) may be set with Invalid Operation Exception (Invalid Compare) for Compare Ordered instructions
- Invalid Operation Exception (SNaN) may be set with Invalid Operation Exception (Invalid Conversion) for Convert To Fixed instructions.
When an exception occurs the instruction execution may be completed or partially completed, depending on the exception and the operation.
For all instructions, except for the Compare and Test instructions, the following exceptions cause the instruction execution to be partially completed. That is, setting of CR field 1 (when Rc=1) and exception status flags is performed, but no result is stored into the target FPR or FPR pair. For Compare and Test instructions, instruction execution is always completed, regardless of whether any DFP exception occurs or not, and whether the exception is enabled or not.
- Enabled Invalid Operation
- Enabled Zero Divide

For the remaining kinds of exceptions, instruction execution is completed, a result, if specified by the instruction, is generated and stored into the target FPR or FPR pair, and appropriate status flags are set. The result may be a different value for the enabled and disabled conditions for some of these exceptions. The kinds of exceptions that deliver a result in target FPR are the following:

- Disabled Invalid Operation
- Disabled Zero Divide
- Disabled Overflow
- Disabled Underflow
- Disabled Inexact
- Enabled Overflow
- Enabled Underflow
- Enabled Inexact

Subsequent sections define each of the DFP exceptions and specify the action that is taken when they are detected.

The IEEE standard specifies the handling of exceptional conditions in terms of "traps" and "trap handlers". In this architecture, a FPSCR exception enable bit of 1 causes generation of the result value specified in the IEEE standard for the "trap enabled" case: the expectation is that the exception will be detected by software, which will revise the result. A FPSCR exception enable bit of 0 causes generation of the "default result" value specified for the "trap disabled" (or "no trap occurs" or "trap is not implemented") case: the expectation is that the exception will not be detected by software, which will simply use the default result. The result to be delivered in each case for each exception is described in the sections below.

The IEEE default behavior when an exception occurs is to generate a default value and not to notify software. In this architecture, if the IEEE default behavior when an exception occurs is desired for all exceptions, all FPSCR exception enable bits should be set to zero and Ignore Exceptions Mode (see below) should be used.

In this case the system floating-point enabled exception error handler is not invoked, even if DFP exceptions occur: software can inspect the FPSCR exception bits if necessary, to determine whether exceptions have occurred.

In this architecture, if software is to be notified that a given kind of exception has occurred, the corresponding FPSCR exception enable bit must be set to one and a mode other than Ignore Exceptions Mode must be used. In this case the system floating-point enabled exception error handler is invoked if an enabled DFP exception occurs. The system floating-point enabled exception error handler is also invoked if a Move To FPSCR instruction causes an exception bit and the corresponding enable bit both to be 1 ; the Move To FPSCR instruction is considered to cause the enabled exception.

The FE0 and FE1 bits control whether and how the system floating-point enabled exception error handler is invoked if an enabled DFP exception occurs. The location of these bits and the requirements for altering them are described in Book III, Power ISA Operating Environment Architecture. (The system floating-point enabled exception error handler is never invoked because of a disabled DFP exception.) The effects of the four possible settings of these bits are as follows.

## FE0 FE1 Description

## 00 Ignore Exceptions Mode

DFP exceptions do not cause the system floating-point enabled exception error handler to be invoked.

01 Imprecise Nonrecoverable Mode
The system floating-point enabled exception error handler is invoked at some point at or beyond the instruction that caused the enabled exception. It may not be possible to identify the excepting instruction or the data that caused the exception. Results produced by the excepting instruction may have been used by or may have affected subsequent instructions that are executed before the error handler is invoked.
10 Imprecise Recoverable Mode
The system floating-point enabled exception error handler is invoked at some point at or beyond the instruction that caused the enabled exception. Sufficient information is provided to the error handler that it can identify the excepting instruction and the operands, and correct the result. No results produced by the excepting instruction have been used by or have affected subsequent instructions that are executed before the error handler is invoked.

## FE0 FE1 Description

11 Precise Mode
The system floating-point enabled exception error handler is invoked precisely at the instruction that caused the enabled exception.

In all cases, the question of whether a DFP result is stored, and what value is stored, is governed by the FPSCR exception enable bits, as described in subsequent sections, and is not affected by the value of the FE0 and FE1 bits.

In all cases in which the system floating-point enabled exception error handler is invoked, all instructions before the instruction at which the system floating-point enabled exception error handler is invoked have completed, and no instruction after the instruction at which the system floating-point enabled exception error handler is invoked has begun execution. (Recall that, for the two Imprecise modes, the instruction at which the system floating-point enabled exception error handler is invoked need not be the instruction that caused the exception.) The instruction at which the system float-ing-point enabled exception error handler is invoked has not been executed unless it is the excepting instruction, in which case it has been executed if the exception is not among those listed on page 186 as suppressed.

> Programming Note - In the ignore and both imprecise modes, a Floating-Point Status and Control Register instruction can be used to force any exceptions, due to instructions initiated before the Floating-Point Status and Control Register instruction, to be recorded in the FPSCR. (This forcing is superfluous for Precise Mode.) In either of the Imprecise modes, a Floating-Point Status and Control Register instruction can be used to force any invocations of the system floating-point enabled exception error handler, due to instructions initiated before the Floating-Point Status and Control Register instruction, to occur. (This forcing has no effect in Ignore Exceptions Mode, and is superfluous for Precise Mode.)

In order to obtain the best performance across the widest range of implementations, the programmer should obey the following guidelines.
■ If the IEEE default results are acceptable to the application, Ignore Exceptions Mode should be used with all FPSCR exception enable bits set to zero.

- If the IEEE default results are not acceptable to the application, Imprecise Nonrecoverable Mode should be used, or Imprecise Recoverable Mode if recoverability is needed, with FPSCR exception
enable bits set to one for those exceptions for which the system floating-point enabled exception error handler is to be invoked.

■ Ignore Exceptions Mode should not, in general, be used when any FPSCR exception enable bits are set to one.

- Precise Mode may degrade performance in some implementations, perhaps substantially, and therefore should be used only for debugging and other specialized applications.


### 5.5.10.1 Invalid Operation Exception

## Definition

An Invalid Operation Exception occurs when an operand is invalid for the specified DFP operation. The invalid DFP operations are:

- Any DFP operation on a signaling $\mathrm{NaN}(\mathrm{SNaN})$, except for Test, Round To DFP Short, Convert To DFP Long, Decode DPD To BCD, Extract Biased Exponent, Insert Biased Exponent, Shift Significand Left Immediate, and Shift Significand Right Immediate
- For add or subtract operations, magnitude subtraction of infinities $(+\infty)+(-\infty)$
- Division of infinity by infinity ( $\infty \div \infty$ )

■ Division of zero by zero ( $0 \div 0$ )

- Multiplication of infinity by zero $(\infty \times 0)$
- Ordered comparison involving a NaN (Invalid Compare)
- The Quantize operation detects that the significand associated with the specified target exponent would have more significant digits than the tar-get-format precision
- For the Quantize operation, when one source operand specifies an infinity and the other specifies a finite number
■ The Reround operation detects that the target exponent associated with the specified target significance would be greater than $\mathrm{X}_{\text {max }}$
- The Encode BCD To DPD operation detects an invalid BCD digit or sign code
- The Convert To Fixed operation involving a number too large in magnitude to be represented in the target format, or involving a NaN .


## Programming Note

In addition, an Invalid Operation Exception occurs if software explicitly requests this by executing an mtfsfi, mtfsf, or mtfsb1 instruction that sets FPSCR ${ }_{\text {VXSOFT }}$ to 1 (Software Request). The purpose of FPSCR ${ }_{V X S O F T}$ is to allow software to cause an Invalid Operation Exception for a condition that is not necessarily associated with the execution of a DFP instruction. For example, it might be set by a program that computes a square root, if the source operand is negative.

## Action

The action to be taken depends on the setting of the Invalid Operation Exception Enable bit of the FPSCR.

When Invalid Operation Exception is enabled (FPSCR ${ }_{V E}=1$ ) and Invalid Operation occurs, the following actions are taken:

1. One or two Invalid Operation Exceptions are set:

$$
\begin{aligned}
& \text { FPSCR }_{V X S N A N} \\
& \text { FPSCR }_{V \times I S I} \\
& \text { FPSCR }_{V \times I D I} \\
& \text { FPSCR }_{V \times Z D Z} \\
& \text { FPSCR }_{V \times I M Z} \\
& \text { FPSCR }_{V \times V C} \\
& \text { FPSCR }_{V \times C V I}
\end{aligned}
$$

(if SNaN )
(if $\infty-\infty$ )
(if $\infty \div \infty$ )
(if $0 \div 0$ )
(if $\infty \times 0$ )
(if invalid comp)
(if invalid conversion)
2. If the operation is an arithmetic, quantum-adjustment, conversion, or format,
the target FPR is unchanged,
FPSCR $_{\text {FR FI }}$ are set to zero, and
FPSCR $_{\text {FPRF }}$ is unchanged.
3. If the operation is a compare,

FPSCR $_{\text {FR FI C }}$ are unchanged, and
FPSCR $_{\text {FPCC }}$ is set to reflect unordered.
When Invalid Operation Exception is disabled (FPSCR ${ }_{V E}=0$ ) and Invalid Operation occurs, the following actions are taken:

1. One or two Invalid Operation Exceptions are set:
FPSCR $_{V X S N A N}$
FPSCR $_{V \times I S I}$
FPSCR $_{V \times I D I}$
FPSCR $_{V X Z D Z}$
FPSCR $_{V \times I M Z}$
FPSCR $_{V X V C}$
FPSCR $_{V X C V I}$
(if SNaN )
(if $\infty-\infty$ )
(if $\infty \div \infty$ )
(if $0 \div 0$ )
(if $\infty \times 0$ )
(if invalid comp)
(if invalid conversion)
2. If the operation is an arithmetic, quantum-adjustment, Round to DFP Long, Convert to DFP Extended, or format
the target FPR is set to a Quiet NaN
FPSCR $_{\text {FR FI }}$ are set to zero
FPSCR $_{\text {FPRF }}$ is set to indicate the class of the result (Quiet NaN)
3. If the operation is a Convert To Fixed
the target FPR is set as follows:
FRT is set to the most positive 64-bit binary integer if the operand in FRB is a positive or
$+\infty$, and to the most negative 64-bit binary integer if the operand in FRB is a negative number, $-\infty$, or NaN .
FPSCR $_{\text {FR FI }}$ are set to zero
FPSCR $_{\text {FPRF }}$ is unchanged
4. If the operation is a compare,

FPSCR $_{\text {FR FI } C}$ are unchanged
FPSCR $_{\text {FPCC }}$ is set to reflect unordered

### 5.5.10.2 Zero Divide Exception

## Definition

A Zero Divide Exception occurs when a Divide instruction is executed with a zero divisor value and a finite nonzero dividend value.

## Action

The action to be taken depends on the setting of the Zero Divide Exception Enable bit of the FPSCR.

When Zero Divide Exception is enabled (FPSCR ${ }_{\text {ZE }}=1$ ) and Zero Divide occurs, the following actions are taken:

1. Zero Divide Exception is set

$$
\mathrm{FPSCR}_{\mathrm{Zx}} \leftarrow 1
$$

2. The target FPR is unchanged
3. FPSCR $_{\text {FR FI }}$ are set to zero
4. FPSCR $_{\text {FPRF }}$ is unchanged

When Zero Divide Exception is disabled (FPSCR ${ }_{\text {ZE }}=0$ ) and Zero Divide occurs, the following actions are taken:

1. Zero Divide Exception is set

$$
\mathrm{FPSCR}_{\mathrm{Zx}} \leftarrow 1
$$

2. The target FPR is set to $\pm \infty$, where the sign is determined by the XOR of the signs of the operands
3. FPSCR $_{\text {FR FI }}$ are set to zero
4. FPSCR $_{\text {FPRF }}$ is set to indicate the class and sign of the result $( \pm \infty)$

### 5.5.10.3 Overflow Exception

## Definition

An overflow exception occurs whenever the target format's largest finite number is exceeded in magnitude by what would have been the rounded result if the exponent range were unbounded.

## Action

Except for Reround, the following describes the handling of the IEEE overflow exception condition. The Reround operation does not recognize an overflow exception condition.
The action to be taken depends on the setting of the Overflow Exception Enable bit of the FPSCR.

When Overflow Exception is enabled (FPSCR ${ }_{\text {OE }}=1$ ) and overflow occurs, the following actions are taken:

1. Overflow Exception is set

FPSCR $_{0 x} \leftarrow 1$
2. The infinitely precise result is divided by $10^{\alpha}$. That is, the exponent adjustment $\alpha$ is subtracted from the exponent. This is called the wrapped result. The exponent adjustment for all operations, except for Round To DFP Short and Round To DFP Long, is 576 for DFP Long and 9216 for DFP Extended. For Round To DFP Short and Round To DFP Long, the exponent adjustment is 192 for the source format of DFP Long and 3072 for the source format of DFP Extended.
3. The wrapped result is rounded to the target-format precision. This is called the wrapped rounded result.
4. If the wrapped rounded result has only one form, it is the delivered result. If the wrapped rounded result has redundant forms and is exact, the result of the form that has the exponent closest to the wrapped ideal exponent is returned. If the wrapped rounded result has redundant forms and is inexact, the result of the form that has the smallest exponent is returned. The wrapped ideal exponent is the result of subtracting the exponent adjustment from the ideal exponent.
5. FPSCR $_{\text {FPRF }}$ is set to indicate the class and sign of the result ( $\pm$ Normal Number)

When Overflow Exception is disabled ( FPSCR $_{\text {OE }}=0$ ) and overflow occurs, the following actions are taken:

1. Overflow Exception is set

FPSCR $_{\text {OX }} \leftarrow 1$
2. Inexact Exception is set

FPSCR $_{X X} \leftarrow 1$
3. The result is determined by the rounding mode and the sign of the intermediate result as follows.

| Rounding Mode | Sign of inter- <br> mediate result |  |
| :--- | :---: | :---: |
|  | Plus | Minus |
|  | $+\infty$ | $-\infty$ |
| Round toward 0 | $+\mathrm{N}_{\max }$ | $-\mathrm{N}_{\max }$ |
| Round toward $+\infty$ | $+\infty$ | $-\mathrm{N}_{\max }$ |
| Round toward - $\infty$ | $+\mathrm{N}_{\max }$ | $-\infty$ |
| Round to Nearest, Ties away <br> from 0 | $+\infty$ | $-\infty$ |
| Round to Nearest, Ties toward 0 | $+\infty$ | $-\infty$ |
| Round away from 0 | $+\infty$ | $-\infty$ |
| Round to prepare for shorter pre- <br> cision | $+\mathrm{N}_{\max }$ | $-\mathrm{N}_{\max }$ |

Figure 78. Overflow Results When Exception Is Disabled
4. The result is placed into the target FPR
5. FPSCR $_{\text {FR }}$ is set to one if the returned result is $\pm \infty$, and is set to zero if the returned result is $\pm N_{\text {max }}$
6. $\mathrm{FPSCR}_{\mathrm{FI}}$ is set to one
7. FPSCR $_{\text {FPRF }}$ is set to indicate the class and sign of the result ( $\pm \infty$ or $\pm$ Normal number)

### 5.5.10.4 Underflow Exception

## Definition

Except for Reround, the following describes the handling of the IEEE underflow exception condition. The Reround operation does not recognize an underflow exception condition.

The Underflow Exception is defined differently for the enabled and disabled states. However, a tininess condition is recognized in both states when a result computed as though both the precision and exponent range were unbounded would be nonzero and less than the target format's smallest normal number, $\mathrm{N}_{\text {min, }}$, in magnitude.

Unless otherwise defined in the instruction description, an underflow exception occurs as follows:

- Enabled:

When the tininess condition is recognized.

- Disabled:

When the tininess condition is recognized and when the delivered result value differs from what would have been computed were both the precision and the exponent range unbounded.

## Action

The action to be taken depends on the setting of the Underflow Exception Enable bit of the FPSCR.
When Underflow Exception is enabled (FPSCR ${ }_{U E}=1$ ) and underflow occurs, the following actions are taken:

1. Underflow Exception is set

$$
\text { FPSCR }_{U X} \leftarrow 1
$$

2. The infinitely precise result is multiplied by $10^{\alpha}$. That is, the exponent adjustment $\alpha$ is added to the exponent. This is called the wrapped result. The exponent adjustment for all operations, except for Round To DFP Short and Round To DFP Long, is 576 for DFP Long and 9216 for DFP Extended. For Round To DFP Short and Round To DFP Long, the exponent adjustment is 192 for the source format of DFP Long and 3072 for the source format of DFP Extended.
3. The wrapped result is rounded to the target-format precision. This is called the wrapped rounded result.
4. If the wrapped rounded result has only one form, it is the delivered result. If the wrapped rounded result has redundant forms and is exact, the result of the form that has the exponent closest to the
wrapped ideal exponent is returned. If the wrapped rounded result has redundant forms and is inexact, the result of the form that has the smallest exponent is returned. The wrapped ideal exponent is the result of adding the exponent adjustment to the ideal exponent.
5. FPSCR FPRF is set to indicate the class and sign of the result ( $\pm$ Normal number)

When Underflow Exception is disabled (FPSCR ${ }_{U E}=0$ ) and underflow occurs, the following actions are taken:

1. Underflow Exception is set

$$
\mathrm{FPSCR}_{U X} \leftarrow 1
$$

2. The infinitely precise result is rounded to the tar-get-format precision.
3. The rounded result is returned. If this result has redundant forms, the result of the form that is closest to the ideal exponent is returned.
4. FPSCR $_{\text {FPRF }}$ is set to indicate the class and sign of the result ( $\pm$ Normal number, $\pm$ Subnormal Number, or $\pm$ Zero)

### 5.5.10.5 Inexact Exception

## Definition

Except for Round to FP Integer Without Inexact, the following describes the handling of the IEEE inexact exception condition. The Round to FP Integer Without Inexact does not recognize an inexact exception condition.
An Inexact Exception occurs when either of two conditions occur during rounding:

1. The delivered result differs from what would have been computed were both the precision and exponent range unbounded.
2. The rounded result overflows and Overflow Exception is disabled.

## Action

The action to be taken does not depend on the setting of the Inexact Exception Enable bit of the FPSCR.

When Inexact Exception occurs, the following actions are taken:

1. Inexact Exception is set

$$
\text { FPSCR }_{X x} \leftarrow 1
$$

2. The rounded or overflowed result is placed into the target FPR
3. FPSCR $_{\text {FPRF }}$ is set to indicate the class and sign of the result

## Programming Note

In some implementations, enabling Inexact Exceptions may degrade performance more than does enabling other types of floating-point exception.

### 5.5.11 Summary of Normal Rounding And Range Actions

Figure 79 and Figure 80 summarize rounding and range actions, with the following exceptions:
■ The Reround operation recognizes neither an underflow nor an overflow exception.

- The Round to FP Integer Without Inexact operation does not recognize the inexact operation exception.

| Range of v | Case | Result (r) <br> when Rounding Mode Is |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | RNE | RNTZ | RNAZ | RAFZ | RTMI | RFSP | RTPI | RTZ |
| v <-Nmax, $\mathrm{q}<-$ Nmax | Overflow | $-\infty^{1}$ | $-\infty^{1}$ | $-\infty^{1}$ | $-\infty^{1}$ | $-\infty^{1}$ | -Nmax | -Nmax | -Nmax |
| $\mathrm{v}<-$ Nmax, $\mathrm{q}=-\mathrm{Nmax}$ | Normal | -Nmax | -Nmax | -Nmax | - | - | -Nmax | -Nmax | -Nmax |
| - Nmax $\leq \mathrm{v} \leq-$ Nmin | Normal | b | b | b | b | b | b | b | b |
| - Nmin $<\mathrm{v} \leq-$ Dmin | Tiny | $\mathrm{b}^{*}$ | $\mathrm{b}^{*}$ | $\mathrm{b}^{*}$ | $\mathrm{b}^{*}$ | $\mathrm{b}^{*}$ | $\mathrm{b}^{*}$ | b | b |
| - Dmin $<\mathrm{v}<-\mathrm{Dmin} / 2$ | Tiny | -Dmin | -Dmin | -Dmin | -Dmin | -Dmin | -Dmin | -0 | -0 |
| $v=-$ Dmin/2 | Tiny | -0 | -0 | -Dmin | -Dmin | -Dmin | -Dmin | -0 | -0 |
| -Dmin/2 < v < 0 | Tiny | -0 | -0 | -0 | -Dmin | -Dmin | -Dmin | -0 | -0 |
| $v=0$ | EZD | +0 | +0 | +0 | +0 | -0 | +0 | +0 | +0 |
| $0<\mathrm{v}<+$ Dmin/2 | Tiny | +0 | +0 | +0 | +Dmin | +0 | +Dmin | +Dmin | +0 |
| $\mathrm{v}=+\mathrm{Dmin} / 2$ | Tiny | +0 | +0 | +Dmin | +Dmin | +0 | +Dmin | +Dmin | +0 |
| +Dmin/2 < v < + Dmin | Tiny | +Dmin | +Dmin | +Dmin | +Dmin | +0 | +Dmin | +Dmin | +0 |
| +Dmin $\leq \mathrm{v}<+$ Nmin | Tiny | $\mathrm{b}^{*}$ | $\mathrm{b}^{*}$ | $\mathrm{b}^{*}$ | $\mathrm{b}^{*}$ | b | $\mathrm{b}^{*}$ | $\mathrm{b}^{*}$ | b |
| +Nmin $\leq \mathrm{v} \leq+$ Nmax | Normal | b | b | b | b | b | b | b | b |
| +Nmax < v, q = +Nmax | Normal | +Nmax | +Nmax | +Nmax | - | +Nmax | +Nmax | - | +Nmax |
| +Nmax < v, q > +Nmax | Overflow | $+\infty^{1}$ | $+\infty^{1}$ | $+\infty^{1}$ | $+\infty^{1}$ | +Nmax | +Nmax | $+\infty^{1}$ | +Nmax |

Explanation:

- This situation cannot occur.

1 The normal result $r$ is considered to have been incremented.

* The rounded value, in the extreme case, may be Nmin. In this case, the exception conditions are underflow, inexact, and incremented.
b The value derived when the precise result v is rounded to the destination's precision, including both bounded precision and bounded exponent range.
$\mathrm{q} \quad$ The value derived when the precise result v is rounded to the destination's precision, but assuming an unbounded exponent range.
$r \quad$ This is the returned value when neither overflow nor underflow is enabled.
v Precise result before rounding, assuming unbounded precision and an unbounded exponent range. For data-format conversion operations, $v$ is the source value.
Dmin Smallest (in magnitude) representable subnormal number in the target format.
EZD The result $r$ of the exact-zero-difference case applies only to ADD and SUBTRACT with both source operands having opposite signs. (For ADD and SUBTRACT, when both source operands have the same sign, the sign of the zero result is the same sign as the sign of the source operands.)
Nmax Largest (in magnitude) representable finite number in the target format.
Nmin $\quad$ Smallest (in magnitude) representable normalized number in the target format.
RAFZ Round away from 0.
RFSP Round to Prepare for Shorter Precision.
RNAZ Round to Nearest, Ties away from 0.
RNE Round to Nearest, Ties to even.
RNTZ Round to Nearest, Ties toward 0.
RTPI Round toward $+\infty$.
RTMI Round toward $-\infty$.
RTZ Round toward 0.
Figure 79. Rounding and Range Actions (Part 1)

| Case | Is $r$ inexact ( $\mathrm{r} \neq \mathrm{v}$ ) | $\mathrm{OE}=1$ | UE=1 | XE=1 | $\begin{gathered} \text { Is r Incre- } \\ \text { mented } \\ (\|\mathrm{rr}\|>\|\mathrm{v}\|) \end{gathered}$ | $\begin{array}{\|c\|} \hline \text { Is } q \\ \text { inexact } \\ (q \neq v) \end{array}$ | $\begin{array}{\|l} \hline \text { Is q Incre- } \\ \text { mented } \\ (\|\mathrm{q}\|>\|\mathrm{v}\|) \end{array}$ | Returned Results and Status Setting* |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Overflow | Yes ${ }^{1}$ | No | - | No | No | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{OX} \leftarrow 1, \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1$ |
| Overflow | Yes ${ }^{1}$ | No | - | No | Yes | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{OX} \mathrm{\leftarrow 1} ,\mathrm{Fl} \mathrm{\leftarrow 1,FR} \mathrm{\leftarrow 1,XX} \mathrm{\leftarrow 1}$ |
| Overflow | Yes ${ }^{1}$ | No | - | Yes | No | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{OX} \leftarrow 1, \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| Overflow | Yes ${ }^{1}$ | No | - | Yes | Yes | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{OX} \leftarrow 1, \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| Overflow | Yes ${ }^{1}$ | Yes | - | - | - | No | No ${ }^{1}$ | $\mathrm{Tw}(\mathrm{q} \div \beta), \mathrm{OX} \leftarrow 1, \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{TO}$ |
| Overflow | Yes ${ }^{1}$ | Yes | - | - | - | Yes | No | $\mathrm{Tw}(\mathrm{q} \div \beta), \mathrm{OX} \leftarrow 1, \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1, \mathrm{TO}$ |
| Overflow | Yes ${ }^{1}$ | Yes | - | - | - | Yes | Yes | $\mathrm{Tw}(\mathrm{q} \div \beta), \mathrm{OX} \leftarrow 1, \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1, \mathrm{TO}$ |
| Normal | No | - | - | - | - | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| Normal | Yes | - | - | No | No | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1$ |
| Normal | Yes | - | - | No | Yes | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1$ |
| Normal | Yes | - | - | Yes | No | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| Normal | Yes | - | - | Yes | Yes | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| Tiny | No | - | No | - | - | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| Tiny | No | - | Yes | - | - | No ${ }^{1}$ | No ${ }^{1}$ | $\mathrm{Tw}(\mathrm{q} \bullet \beta), \mathrm{UX} \leftarrow 1, \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{TU}$ |
| Tiny | Yes | - | No | No | No | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{UX} \leftarrow 1, \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1$ |
| Tiny | Yes | - | No | No | Yes | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{UX} \leftarrow 1, \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1$ |
| Tiny | Yes | - | No | Yes | No | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{UX} \leftarrow 1, \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| Tiny | Yes | - | No | Yes | Yes | - | - | $\mathrm{T}(\mathrm{r}), \mathrm{UX} \leftarrow 1, \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| Tiny | Yes | - | Yes | - | - | No | No ${ }^{1}$ | $\mathrm{Tw}(\mathrm{q} \bullet \beta), \mathrm{UX} \leftarrow 1, \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{TU}$ |
| Tiny | Yes | - | Yes | - | - | Yes | No | $\mathrm{Tw}(\mathrm{q} \bullet \beta), \mathrm{UX} \leftarrow 1, \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1, \mathrm{TU}$ |
| Tiny | Yes | - | Yes | - | - | Yes | Yes | $\mathrm{Tw}(\mathrm{q} \bullet \beta), \mathrm{UX} \leftarrow 1, \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1, \mathrm{TU}$ |

Explanation:

- The results do not depend on this condition.

1 This condition is true by virtue of the state of some condition to the left of this column.

* Rounding sets only the FI and FR status flags. Setting of the OX, XX, or UX flag is part of the exception actions. They are listed here for reference.
$\beta \quad$ Wrap adjust, which depends on the type of operation and operand format. For all operations except Round to DFP Short and Round to DFP Long, the wrap adjust depends on the target format: $\beta=10^{\alpha}$, where $\alpha$ is 576 for DFP Long, and 9216 for DFP Extended. For Round to DFP Short and Round to DFP Long, the wrap adjust depends on the source format: $\beta=10^{\kappa}$ where $\kappa$ is 192 for DFP Long and 3072 for DFP Extended.
q The value derived when the precise result v is rounded to destination's precision, but assuming an unbounded exponent range.
$r$ The result as defined in Part 1 of this figure.
v Precise result before rounding, assuming unbounded precision and unbounded exponent range.
FI Floating-Point-Fraction-Inexact status flag, FPSCR $_{\text {FI }}$. This status flag is non-sticky.
FR Floating-Point-Fraction-Rounded status flag, FPSCR $_{\text {FR }}$.
OX Floating-Point Overflow Exception status flag, $\mathrm{FPSCR}_{0} \mathrm{x}$.
TO The system floating-point enabled exception error handler is invoked for the overflow exception if the FE0 and FE1 bits in the machine-state register are set to any mode other than the ignore-exception mode.
TU The system floating-point enabled exception error handler is invoked for the underflow exception if the FE0 and FE1 bits in the machine-state register are set to any mode other than the ignore-exception mode.
TX The system floating-point enabled exception error handler is invoked for the inexact exception if the FE0 and FE1 bits in the machine-state register are set to any mode other than the ignore-exception mode.
$\mathrm{T}(\mathrm{x}) \quad$ The value x is placed at the target operand location.
$\mathrm{Tw}(\mathrm{x}) \quad$ The wrapped rounded result x is placed at the target operand location. For all operations except data format conversions, the wrapped rounded result is in the same format and length as normal results at the target location. For data format conversions, the wrapped rounded result is in the same format and length as the source, but rounded to the target-format precision.
UX Floating-Point-Underflow-Exception status flag, FPSCR $_{U X}$
XX Float-Point-Inexact-Exception Status flag, FPSCR XXX $^{\text {. The flag is a sticky version of } \text { FPSCR }_{\text {FI }} \text {. When FPSCR }}$ FI is set to a new value, the new value of FPSCR $_{X X}$ is set to the result of ORing the old value of FPSCR $_{X X}$ with the new value of FPSCR $_{\text {FI }}$.
Figure 80. Rounding and Range Actions (Part 2)


### 5.6 DFP Instruction Descriptions

The following sections describe the DFP instructions. When a 128-bit operand is used, it is held in a FPR pair and the instruction mnemonic uses a letter " $q$ " to mean the quad-precision operation. Note that in the following descriptions, FPXp denotes a FPR pair and must address an even-odd pair. If the FPXp field specifies an odd-numbered register, then the instruction form is invalid. The notation FPX[p] means either a FPR, FPX, or a FPR pair, FPXp.
For DFP instructions, if a DFP operand is returned, the trailing significand field of the target operand is encoded using preferred DPD codes.

### 5.6.1 DFP Arithmetic Instructions

All DFP arithmetic instructions are X-form instructions. They all set the FI and FR status flags, and also set the FPSCR $_{\text {FPRF }}$ field. Furthermore, they all have an ideal exponent assigned and employ the record bit (Rc).

The arithmetic instructions consist of Add, Divide, Multiply, and Subtract.

| DFP Add [Quad] |  |  |  |  | $\begin{array}{r} X \text {-form } \\ (R c=0) \\ (R c=1) \end{array}$ |  | DFP Subtract [Quad] |  |  |  |  | X-form |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| dadd <br> dadd. | FRT,FRA,FRB FRT,FRA,FRB |  |  |  |  |  | dsub dsub. | FRT,FRA,FRB FRT,FRA,FRB |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| $59$ | ${ }_{6}$ FRT | $\begin{aligned} & \text { FRA } \\ & 11 \end{aligned}$ | $\begin{aligned} & \text { FRB } \\ & 16 \end{aligned}$ |  | 2 | Rc <br> 31 | 059 | ${ }_{6}$ FRT | 11. | ${ }_{16} \mathrm{FRB}$ | 21 | 514 | Rc <br> 31 |
| daddq daddq. | $\begin{aligned} & \text { FRT } \\ & \text { FRT } \end{aligned}$ | ,FRAp, <br> ,FRAp, | $\begin{aligned} & \text { FRBp } \\ & \text { FRBp } \end{aligned}$ |  |  | $\begin{aligned} & (R c=0) \\ & (R c=1) \end{aligned}$ | dsubq dsubq | FRT | $p, F R A p$, <br> ,FRAp, | $\begin{aligned} & \text { FRBp } \\ & \text { FRBp } \end{aligned}$ |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |
| $\begin{array}{\|r} \hline \end{array} 63$ | $\begin{aligned} & \text { FRTp } \\ & 6 \\ & \hline \end{aligned}$ | $\begin{array}{\|l} \hline \text { FRAp } \\ 11 \\ \hline \end{array}$ | $\begin{aligned} & \text { FRBp } \\ & 16 \\ & \hline \end{aligned}$ | 21 | 2 | Rc <br> 31 | 063 | $\left.\right\|_{6} \mathrm{FRTp}$ | FRAp <br> 11 | FRBp <br> 16 | 21 | 514 | Rc <br> 31 |

The DFP operand in FRA[p] is added to the DFP operand in FRB[p].

The result is rounded to the target-format precision under control of the DRN (bits 29:31) of the FPSCR. An appropriate form of the rounded result is selected based on the ideal exponent and is placed in FRT[p]. The ideal exponent is the smaller exponent of the two source operands.

Figure 81 summarizes the actions for Add. Figure 81 does not include the setting of the FPSCR FPRF field. The FPSCR FPRRF field is always set to the class and sign of the result, except for an enabled invalid-operation exception, in which case the field remains unchanged.

## Special Registers Altered:

FPRF FR FI
FX OX UX XX
VXSNAN VXISI
CR1
(if $\mathrm{Rc}=1$ )
dsubq

The DFP operand in FRB[p] is subtracted from the DFP operand in FRA[p].

The result is rounded to the target-format precision under control of the DRN (bits 29:31) of the FPSCR. An appropriate form of the rounded result is selected based on the ideal exponent and is placed in FRT[p]. The ideal exponent is the smaller exponent of the two source operands.
The execution of Subtract is identical to that of Add, except that the operand in FRB participates in the operation with its sign bit inverted. See Figure 81. The table does not include the setting of the FPSCR FPRF field. The FPSCR FPRF field is always set to the class and sign of the result, except for an enabled invalid-operation exception, in which case the field remains unchanged.

```
Special Registers Altered:
    FPRF FR FI
    FX OX UX XX
    VXSNAN VXISI
    CR1
    (if Rc=1)
```


## Version 3.0

| Operand a in FRA[p] is | Actions for Add ( $a+b$ ) when operand b in FRB[p] is |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | - $\infty$ | F | $+\infty$ | QNaN | SNaN |
| - $\infty$ | T(-dINF) | T(-dINF) | $\mathrm{V}_{\mathrm{XISI}}: \mathrm{T}(\mathrm{dNaN})$ | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| F | T(-dINF) | $S(a+b)$ | T(+dINF) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| + | $\mathrm{V}_{\mathrm{XISI}}$ : T(dNaN) | T(+dINF) | T(+dINF) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| QNaN | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| SNaN | $\mathrm{V}_{\text {XSNAN: }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(a) | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ |
| Explanation: |  |  |  |  |  |
| $a+b$ | The value $a$ added to $b$, rounded to the target-format precision and returned in the appropriate form. (See Section 5.5.11 on page 192) |  |  |  |  |
| +dINF | Default plus infinity. |  |  |  |  |
| - dINF | Default minus infinity. |  |  |  |  |
| dNaN | Default quiet NaN . |  |  |  |  |
| F | All finite numbers, including zeros. |  |  |  |  |
| $\mathrm{P}(\mathrm{x})$ | The QNaN of operand x is propagated and placed in FRT[p]. |  |  |  |  |
| $S(x)$ | The value $x$ is placed in FRT[p] with the sign set by the rules of algebra. When the source operands have the same sign, the sign of the result is the same as the sign of the operands, including the case when the result is zero. When the operands have opposite signs, the sign of a zero result is positive in all rounding modes, except round toward $-\infty$, in which case, the sign is minus. |  |  |  |  |
| T(x) | The value x is placed in FRT[p]. |  |  |  |  |
| $\mathrm{U}(\mathrm{x})$ | The SNaN of operand x is converted to the corresponding QNaN and placed in FRT[p]. |  |  |  |  |
| $\mathrm{V}_{\mathrm{XISI}}$ | The Invalid-Operation Exception (VXISI) occurs. The result is produced only when the exception is disabled. (See Section 5.5.10.1 "Invalid Operation Exception" on page 188 for the exception actions.) |  |  |  |  |
| $\mathrm{V}_{\text {XSNAN }}$ | The Invalid-Operation Exception (VXSNAN) occurs. The result is produced only when the exception is disabled. (See Section 5.5.10.1 "Invalid Operation Exception" on page 188 for the exception actions.) |  |  |  |  |

Figure 81. Actions: Add
DFP Multiply [Quad]
X-form


| 63 | FRTp | FRAp | FRBp |  | 34 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

The DFP operand in FRA[p] is multiplied by the DFP operand in FRB[p].

The result is rounded to the target-format precision under control of the DRN (bits 29:31) of the FPSCR. An appropriate form of the rounded result is selected based on the ideal exponent and is placed in FRT[p]. The ideal exponent is the sum of the two exponents of the source operands.

Figure 82 summarizes the actions for Multiply. Figure 82 does not include the setting of the FPSCR $_{\text {F- }}$ PRF field. The FPSCR ${ }_{\text {FPRF }}$ field is always set to the class and sign of the result, except for an enabled
invalid-operation exception, in which case the field remains unchanged.

## Special Registers Altered:

FPRF FR FI
FX OX UX XX VXSNAN VXIMZ
CR1
(if $\mathrm{Rc}=1$ )

| Operand a in FRA[p] is | Actions for Multiply (a*b) when operand b in FRB[p] is |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0 | Fn | $\infty$ | QNaN | SNaN |
| 0 | S(a*b) | $S\left(a^{*}{ }^{\text {b }}\right.$ | $\mathrm{V}_{\mathrm{XIMZ}}$ : $\mathrm{T}(\mathrm{dNaN})$ | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| Fn | S(a* b) | $S\left({ }^{*}{ }^{\text {b }}\right.$ ) | S(dINF) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| $\infty$ | $\mathrm{V}_{\text {XIMZ }}$ : T(dNaN) | S(dINF) | S(dINF) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| QNaN | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| SNaN | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ |

Explanation:
$a$ * $b \quad$ The value a multiplied by $b$, rounded to the target-format precision and returned in the appropriate form. (See Section 5.5.11 on page 192)
dINF Default infinity.
dNaN Default quiet NaN.
Fn Finite nonzero number (includes both normal and subnormal numbers).
$P(x) \quad$ The QNaN of operand $x$ is propagated and placed in FRT[p].
$S(x) \quad$ The value $x$ is placed in FRT[p] with the sign set to the exclusive-OR of the source-operand signs.
$\mathrm{T}(\mathrm{x}) \quad$ The value x is placed in FRT[p].
$\mathrm{U}(\mathrm{x}) \quad$ The SNaN of operand x is converted to the corresponding QNaN and placed in FRT[p].
$\mathrm{V}_{\mathrm{XIMZ}}: \quad$ The Invalid-Operation Exception (VXIMZ) occurs. The result is produced only when the exception is disabled. (See Section 5.5.10.1 "Invalid Operation Exception" on page 188 for the exception actions.)
$V_{\text {XSNAN: }} \quad$ The Invalid-Operation Exception (VXSNAN) occurs. The result is produced only when the exception is disabled. (See Section 5.5.10.1 "Invalid Operation Exception" on page 188 for the exception actions.)

Figure 82. Actions: Multiply

## DFP Divide [Quad]

X-form

| ddiv ddiv. | FRT,FRA,FRB FRT,FRA,FRB |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $59$ | ${ }_{6} \text { FRT }$ | $\begin{aligned} & \text { FRA } \\ & 11 \end{aligned}$ | $\begin{aligned} & \text { FRB } \\ & 16 \end{aligned}$ | $21546$ | Rc 31 |
| ddivq <br> ddivq. | FRTp,FRAp,FRBp |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |


| 63 | FRTp | FRAp | FRBp |  | 546 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 |  |
| 0 |  |  |  |  |  |

The DFP operand in FRA[p] is divided by the DFP operand in FRB[p].

The result is rounded to the target-format precision under control of the DRN (bits 29:31) of the FPSCR. An appropriate form of the rounded result is selected based on the ideal exponent and is placed in FRT[p]. The ideal exponent is the difference of subtracting the exponent of the divisor from the exponent of the dividend.

Figure 83 summarizes the actions for Divide. Figure 83 does not include the setting of the FPSCR FPRF $^{\text {field. }}$ The FPSCR FPRF field is always set to the class and sign of the result, except for an enabled invalid-operation and enabled zero-divide exceptions, in which cases the field remains unchanged.

## Special Registers Altered:

FPRF FR FI
FX OX UX ZX XX
VXSNAN VXIDI VXZDZ
CR1
(if $\mathrm{Rc}=1$ )

| Operand a in FRA[p] is | Actions for Divide (a $\div$ b) when operand b in FRB[p] |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0 | Fn | $\infty$ | QNaN | SNaN |
| 0 | $\mathrm{V}_{\text {XZDZ }}:$ T(dNaN) | $S(a \div b)$ | S(zt) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| Fn | Zx: S(dINF) | $S(a \div b)$ | S(zt) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| $\infty$ | S(dINF) | S(dINF) | $\mathrm{V}_{\text {XIDI }}: \mathrm{T}(\mathrm{dNaN})$ | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| QNaN | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| SNaN | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN: }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ |
| Explanation: |  |  |  |  |  |
| $a \div b$ | The value a divided by b, rounded to the target-format precision and returned in the appropriate form. (See Section 5.5 .11 on page 192.) |  |  |  |  |
| dINF | Default infinity. |  |  |  |  |
| dNaN | Default quiet NaN . |  |  |  |  |
| Fn | Finite nonzero number (includes both normal and subnormal numbers). |  |  |  |  |
| $\mathrm{P}(\mathrm{x})$ | The QNaN of operand $x$ is propagated and placed in FRT[p]. |  |  |  |  |
| $S(x)$ | The value $x$ is placed in FRT[p] with the sign set to the exclusive-OR of the source-operand signs. |  |  |  |  |
| T(x) |  |  |  |  |  |
| $\mathrm{U}(\mathrm{x})$ | The SNaN of operand x is converted to the corresponding QNaN and placed in FRT[p]. |  |  |  |  |
| $\mathrm{V}_{\text {XIDI: }}$ : | The Invalid-Operation Exception (VXIDI) occurs. The result is produced only when the exception is disabled. (See Section 5.5.10.1 "Invalid Operation Exception" on page 188 for the exception actions.) |  |  |  |  |
| $\mathrm{V}_{\text {XSNAN: }}$ | The Invalid-Operation Exception (VXSNAN) occurs. The result is produced only when the exception is disabled. (See Section 5.5.10.1 "Invalid Operation Exception" on page 188 for the exception actions.) |  |  |  |  |
| $\mathrm{V}_{\mathrm{XZDZ}}$ | The Invalid-Operation Exception (VXZDZ) occurs. The result is produced only when the exception is disabled. (See Section 5.5.10.1 "Invalid Operation Exception" on page 188 for the exception actions.) |  |  |  |  |
| zt | True zero (zero significand and most negative exponent). |  |  |  |  |
| Zx | The Zero-Divide Exception occurs. The result is produced only when the exception is disabled (See Section 5.5.10.2 "Zero Divide Exception" on page 189 for the exception actions.) |  |  |  |  |

Figure 83. Actions: Divide

### 5.6.2 DFP Compare Instructions

The DFP compare instructions consist of the Compare Ordered and Compare Unordered instructions. The compare instructions do not provide the record bit.

The comparison sets the designated CR field to indicate the result. The FPSCR ${ }_{\text {FPCC }}$ is set in the same way.

The codes in the CR field BF and FPSCR FPCC are defined for the DFP compare operations as follows.

Bit Name Description
0 FL $\quad($ FRA $[p])<(F R B[p])$
1 FG (FRA[p]) $>($ FRB[p])
$2 \mathrm{FE} \quad(\mathrm{FRA}[\mathrm{p}])=(\mathrm{FRB}[\mathrm{p}])$
3 FU (FRA[p]) ? (FRB[p])

## Version 3.0

## DFP Compare Unordered [Quad] X-form

dcmpu BF,FRA,FRB

| 59 | BF | $/ /$ | FRA | FRB |  | 642 | $/$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 9 | 11 | 16 | 21 |  | 31 |

dcmpuq BF,FRAp,FRBp

| 63 | BF | $/ /$ | FRAp | FRBp |  | 642 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 9 | 11 | 16 | 21 |  |

The DFP operand in FRA[p] is compared to the DFP operand in $\operatorname{FRB}[p]$. The result of the compare is placed into CR field BF and the FPSCR FPCC .

## Special Registers Altered:

CR field BF
FPCC
FX VXSNAN

| Operand a in FRA[p] is | Actions for Compare Unordered (a:b) when operand b in FRB[p] is |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | $-\infty$ | F | $+\infty$ | QNaN | SNaN |
| $-\infty$ | AeqB | AltB | AltB | AuoB | Fu, $\mathrm{V}_{\text {XSNAN }}$ |
| F | AgtB | C(a:b) | AltB | AuoB | Fu, $\mathrm{V}_{\text {XSNAN }}$ |
| $+\infty$ | AgtB | AgtB | AeqB | AuoB | Fu, $\mathrm{V}_{\text {XSNAN }}$ |
| QNaN | AuoB | AuoB | AuoB | AuoB | Fu, $\mathrm{V}_{\text {XSNAN }}$ |
| SNaN | Fu, $\mathrm{V}_{\text {XSNAN }}$ | Fu, $\mathrm{V}_{\text {XSNAN }}$ | Fu, $\mathrm{V}_{\text {XSNAN }}$ | Fu, $\mathrm{V}_{\text {XSNAN }}$ | Fu, $\mathrm{V}_{\text {XSNAN }}$ |
| Explanation: |  |  |  |  |  |
| $C(a: b)$ | Algebraic comparison. See the table below. |  |  |  |  |
| F | All finite numbers, including zeros. |  |  |  |  |
| AeqB | CR field BF and FPSCR FPCC $^{\text {are set to } 0 \mathrm{~b} 0010}$. |  |  |  |  |
| AgtB | CR field BF and FPSCR FPCC $^{\text {are set to } 0 \mathrm{~b} 0100}$. |  |  |  |  |
| AltB | $C R$ field BF and FPSCR ${ }_{\text {FPCC }}$ are set to 0b1000. |  |  |  |  |
| Auob | CR field BF and FPSCR ${ }_{\text {FPCC }}$ are set to 0b0001. |  |  |  |  |
| $V_{\text {XSNAN }}$ | The invalid-operation exception (VXSNAN) occurs. See Section 5.5.10.1 for actions. |  |  |  |  |


| Relation of Value a to Value b | Action for C(a:b) |
| :---: | :---: |
| $\mathrm{a}=\mathrm{b}$ | AeqB |
| $\mathrm{a}<\mathrm{b}$ | AltB |
| $\mathrm{a}>\mathrm{b}$ | AgtB |

Figure 84. Actions: Compare Unordered

## DFP Compare Ordered [Quad] X-form

dcmpo BF,FRA,FRB

| 59 | BF | // | FRA | FRB |  | 130 | $/$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 9 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |  |

dcmpoq BF,FRAp,FRBp

| 63 | BF | // | FRAp | FRBp |  | 130 | $/$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 9 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |  |

The DFP operand in FRA[p] is compared to the DFP operand in $\operatorname{FRB}[p]$. The result of the compare is placed into CR field BF and the FPSCR FPCC .

Special Registers Altered:
CR field BF
FPCC
FX VXSNAN VXVC

| Operand a in FRA[p] is | Actions for Compare ordered (a:b) when operand $b$ in FRB[p] is |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | - $\infty$ | F | $+\infty$ | QNaN | SNaN |
| $-\infty$ | AeqB | AltB | AltB | AuoB, $\mathrm{V}_{\mathrm{XVc}}$ | AuoB, $\mathrm{V}_{\mathrm{XSV}}$ |
| F | AgtB | $\mathrm{C}(\mathrm{a}: \mathrm{b})$ | AltB | AuoB, $\mathrm{V}_{\mathrm{XVc}}$ | AuoB, $\mathrm{V}_{\mathrm{XSV}}$ |
| $+\infty$ | AgtB | AgtB | AeqB | AuoB, $\mathrm{V}_{\mathrm{XVC}}$ | AuoB, $\mathrm{V}_{\mathrm{XSV}}$ |
| QNaN | AuoB, $\mathrm{V}_{\mathrm{XVC}}$ | AuoB, $\mathrm{V}_{\mathrm{XVC}}$ | AuoB, $\mathrm{V}_{\mathrm{XVC}}$ | AuoB, $\mathrm{V}_{\mathrm{XVC}}$ | AuoB, $\mathrm{V}_{\text {XSV }}$ |
| SNaN | AuoB, $\mathrm{V}_{\mathrm{XSV}}$ | AuoB, $\mathrm{V}_{\mathrm{XSV}}$ | AuoB, $\mathrm{V}_{\mathrm{XSV}}$ | AuoB, $\mathrm{V}_{\mathrm{XSV}}$ | AuoB, $\mathrm{V}_{\mathrm{XSV}}$ |

Explanation:
$\mathrm{C}(\mathrm{a}: \mathrm{b}) \quad$ Algebraic comparison. See the table below
F All finite numbers, including zeros
AeqB $\quad$ CR field BF and FPSCR ${ }_{\text {FPCC }}$ are set to 0b0010.
AgtB $\quad$ CR field BF and FPSCR FPCC are set to 0b0100.
AltB $\quad$ CR field BF and FPSCR FPCC are set to 0b1000.
AuoB CR field BF and FPSCR ${ }_{\text {FPCC }}$ are set to 0b0001.
$V_{X S V} \quad$ The invalid-operation exception (VXSNAN) occurs. Additionally, if the exception is disabled ( FPSCR $_{V E}=0$ ), then FPSCR $_{V X V C}$ is also set to one. See Section 5.5.10.1 for actions.
$\mathrm{V}_{\mathrm{XVC}}$ The invalid-operation exception (VXVC) occurs. See Section 5.5.10.1 for actions.

| Relation of Value a to Value b | Action for C(a:b) |
| :---: | :---: |
| $\mathrm{a}=\mathrm{b}$ | AeqB |
| $\mathrm{a}<\mathrm{b}$ | AltB |
| $\mathrm{a}>\mathrm{b}$ | AgtB |

Figure 85. Actions: Compare Ordered

### 5.6.3 DFP Test Instructions

The DFP test instructions consist of the Test Data Class, Test Data Group, Test Exponent, and Test Significance instructions, and they do not provide the record bit.

The test instructions set the designated CR field to indicate the result. The FPSCR ${ }_{\text {FPCC }}$ is set in the same way.

## DFP Test Data Class [Quad] <br> Z22-form

dtstdc BF,FRA,DCM

| 59 | BF | // | FRA | DCM | 194 | 1 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 9 | 11 | 16 | 22 |  |

dtstdcq
BF,FRAp,DCM

| 63 | BF | I/ | FRAp | DCM | 194 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 9 | 11 | 16 | 22 |

Let the DCM (Data Class Mask) field specify one or more of the 6 possible data classes, where each bit corresponds to a specific data class.

| DCM Bit | Data Class |
| :--- | :--- |
| 0 | Zero |
| 1 | Subnormal |
| 2 | Normal |
| 3 | Infinity |
| 4 | Quiet NaN |
| 5 | Signaling NaN |

CR field BF and FPSCR ${ }_{\text {FPCC }}$ are set to indicate the sign of the DFP operand in FRA[p] and whether the data class of the DFP operand in FRA[p] matches any of the data classes specified by DCM.

| Field | Meaning |
| :--- | :--- |
| 0000 | Operand positive with no match |
| 0010 | Operand positive with match |
| 1000 | Operand negative with no match |
| 1010 | Operand negative with match |
| Special Registers Altered: |  |
| CR field BF |  |
| FPCC |  |

DFP Test Data Group [Quad] Z22-form
dtstdg BF,FRA,DGM

| 59 | BF | $/ /$ | FRA | DGM |  | 226 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 9 | 11 | 16 | 22 |  |

dtstdgq BF,FRAp,DGM

| 63 | BF | $/ /$ | FRAp | DGM |  | 226 | $/$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 9 | 11 | 16 | 22 |  | 31 |

Let the DGM (Data Group Mask) field specify one or more of the 6 possible data groups, where each bit corresponds to a specific data group.
The term extreme exponent means either the maximum exponent, $X_{\text {max }}$, or the minimum exponent, $X_{\text {min }}$.

## DGM Bit Data Group

$0 \quad$ Zero with non-extreme exponent
1 Zero with extreme exponent
2 Subnormal or (Normal with extreme exponent)
3 Normal with non-extreme exponent and leftmost zero digit in significand
4 Normal with non-extreme exponent and leftmost nonzero digit in significand 5 Special symbol (Infinity, QNaN, or SNaN)

CR field BF and FPSCR ${ }_{\text {FPCC }}$ are set to indicate the sign of the DFP operand in FRA[p] and whether the data group of the DFP operand in FRA[p] matches any of the data groups specified by DGM.

| Field | Meaning |
| :--- | :--- |
| 0000 | Operand positive with no match |
| 0010 | Operand positive with match |
| 1000 | Operand negative with no match |
| 1010 | Operand negative with match |

## Special Registers Altered:

CR field BF
FPCC

## DFP Test Exponent [Quad] X-form

dtstex BF,FRA,FRB

| 59 | BF | $/ /$ | FRA | FRB |  | 162 | $/$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 9 | 11 | 16 | 21 |  | 31 |

dtstexq BF,FRAp,FRBp

| 63 | BF | $/ /$ | FRAp | FRBp |  | 162 | $/$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 9 | 11 | 16 | 21 |  | 31 |

The exponent value (Ea) of the DFP operand in FRA[p] is compared to the exponent value (Eb) of the DFP operand in FRB [p]. The result of the compare is placed into CR field BF and the FPSCR ${ }_{\text {FPCC }}$.

The codes in the CR field BF and FPSCR FPCC are defined for the DFP Test Exponent operations as follows.

Bit Description
$0 \quad \mathrm{Ea}<\mathrm{Eb}$
$1 \quad \mathrm{Ea}>\mathrm{Eb}$
$2 \quad E a=E b$
3 Ea? Eb

## Special Registers Altered:

CR field BF
FPCC

| Operand a in |
| :---: | :---: | :---: | :---: | :---: |
| FRA[p] is |$\quad$ Actions for Test Exponent (Ea:Eb) when operand b in FRB[p] is


| Relation of Value Ea to Value Eb | Action for C(Ea:Eb) |
| :---: | :---: |
| $\mathrm{Ea}=\mathrm{Eb}$ | AeqB |
| $\mathrm{Ea}<\mathrm{Eb}$ | AltB |
| $\mathrm{Ea}>\mathrm{Eb}$ | AgtB |

Figure 86. Actions: Test Exponent


Let $k$ be the contents of bits $58: 63$ of FPR[FRA] that specifies the reference significance.

For dtstsf, let the value NSDb be the number of significant digits of the DFP value in FPR[FRB].

For dtstsfq, let the value NSDD be the number of significant digits of the DFP value in FPR[FRBp: FRBp +1$]$.

For this instruction, the number of significant digits of the value 0 is considered to be zero.
$N S D b$ is compared to $k$. The result of the compare is placed into CR field BF and the FPCC as follows.

## Bit Description

$0 \quad k \neq 0$ and $k<N S D b$
$1 \quad k \neq 0$ and $k>$ NSDb, or $k=0$
$2 \quad k \neq 0$ and $k=N S D b$
3 k ? NSDb

## Special Registers Altered:

CR field BF
FPCC

| Actions for Test Significance |  |  |  |
| :---: | :---: | :---: | :---: |
| when the operand in VSR[FRB] or VSR[FRBp:FRBp+1] is |  |  |  |
| F | $\infty$ | QNaN | SNaN |
| $C(U \mid M: N S D b)$ | $A \cup 0 B$ | $A \cup 0 B$ | $A \cup 0 B$ |

Explanation:

| $C(k: N S D b)$ | Algebraic comparison. See the table <br> below. |
| :---: | :--- |
| $F$ | All finite numbers, including zeros. |
| $A e q B$ | $C R$ field $B F$ and $F P C C$ are set to $0 b 0010$. |
| $A g t B$ | $C R$ field $B F$ and $F P C C$ are set to $0 b 0100$. |
| $A I t B$ | $C R$ field BF and FPCC are set to $0 b 1000$. |
| $A u O B$ | $C R$ field $B F$ and FPCC are set to $0 b 0001$. |


| Relation of Value NSDb to Value $\mathbf{k}$ | Action for C(k:NSDb) |
| :--- | :---: |
| $k \neq 0$ and $k=N S D b$ | AeqB |
| $k \neq 0$ and $k<N S D b$ | Al t B |
| $k \neq 0$ and $k>N S D b$, or $k=0$ | Agt B |

Figure 87. Actions: Test Significance

## Programming Note

The reference significance can be loaded into a FPR using a Load Float as Integer Word Algebraic instruction

## DFP Test Significance Immediate [Quad]

dtstsfi BF,UIM,FRB

| 59 | BF | 1 | UIM | FRB |  | 675 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 1 |  |  |  |
| 31 |  |  |  |  |  |  |

dtstsfiq BF,UIM,FRBp

| 63 | ${ }_{6} \mathrm{BF}$ | $l_{9}^{\prime}$ | UIM | FRBp |  | 675 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  |  |  |  |  |  |
| 10 |  |  |  |  |  |  |

Let the value UIM specify the reference significance.
For dtstsfi, let the value $N S D b$ be the number of significant digits of the DFP value in FPR[FRB].

For dtstsfiq, let the value $N S D b$ be the number of significant digits of the DFP value in FPR[FRBp: FRBp+1].

For this instruction, the number of significant digits of the value 0 is considered to be zero.

NSDb is compared to $\mathrm{UI} M$. The result of the compare is placed into CR field BF and the FPCC as follows.

```
    Bit
        Description
O UIM\not=O and UIM<NSDb
1.UIM}\not=0\mathrm{ and UIM }>NNSDb, or UIM=
    UIM }\not=0\mathrm{ and UIM = NSDb
    UIM? NSDb
```

Special Registers Altered:
CR field BF
FPCC

| Actions for Test Significance <br> when the operand in VSR[FRB] or VSR[FRBp:FRBp+1] is <br> F$\infty^{\circ}$ |  |  |  |
| :---: | :---: | :---: | :---: |
| QNaN | SNaN |  |  |
| C(U\|M: NSDb) | AuOB | AuoB | AuoB |

Explanation:

| $C(U \mid M: N S D b)$ | Algebraic comparison. See the table <br> below. |
| :---: | :--- |
| $F$ | All finite numbers, including zeros. |
| $A e q B$ | $C R$ field $B F$ and $F P C C$ are set to $0 b 0010$. |
| $A g t B$ | $C R$ field $B F$ and $F P C C$ are set to $0 b 0100$. |
| $A I t B$ | $C R$ field $B F$ and $F P C C$ are set to $0 b 1000$. |
| $A \cup O B$ | $C R$ field $B F$ and $F P C C$ are set to $0 b 0001$. |


| Relation of Value NSDb to Value UIM | Action for C(UIM:NSDb) |
| :--- | :---: |
| $U I M \neq 0$ and $U I M=N S D b$ | Aeq $B$ |
| $U I M \neq 0$ and $U I M<N S D b$ | AI t B |
| $U I M \neq 0$ and $U I M>N S D b$, or $U I M=0$ | Ag B |

Figure 88. Actions: Test Significance

### 5.6.4 DFP Quantum Adjustment Instructions

The Quantum Adjustment operations consist of the Quantize, Quantize Immediate, Reround, and Round To FP Integer operations.

The Quantum Adjustment instructions are Z23-form instructions and have an immediate RMC (Round-ing-Mode-Control) field, which specifies the rounding mode used. For Quantize, Quantize Immediate, and Reround, the RMC field contains the primary encoding. For Round to FP Integer, the field contains either pri-
mary or secondary encoding, depending on the setting of a RMC-encoding-selection bit. See Section 5.5.2 "Rounding Mode Specification" on page 183 for the definition of RMC encoding.

All Quantum Adjustment instructions set the FI and FR status flags, and also set the FPSCR FPRF field. The record bit is provided to each of these instructions. They return the target operand in a form with the ideal exponent.

DFP Quantize Immediate [Quad]Z23-form

| dquai | TE,FRT,FRB,RMC | $(R \mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| dquai. | TE,FRT,FRB,RMC | $(\mathrm{Rc}=1)$ |


| 59 | FRT | TE | FRB | RMC | 67 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 | 21 | 23 |  |


| dquaiq | TE,FRTp,FRBp,RMC | $(R \mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| dquaiq. | TE,FRTp,FRBp,RMC | $(\mathrm{Rc}=1)$ |


| 63 | FRTp | TE | FRBp | RMC |  | 67 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 | 21 | 23 |  | 31 |

The DFP operand in FRB[p] is converted and rounded to the form with the exponent specified by TE based on the rounding mode specified in the RMC field. TE is a 5 -bit signed binary integer. The result of that form is placed in FRT[p]. The sign of the result is the same as the sign of the operand in FRB[p]. The ideal exponent is the exponent specified by TE.

When the value of the operand in $\mathrm{FRB}[\mathrm{p}]$ is greater than $\left(10^{\mathrm{p}}-1\right) \times 10^{\mathrm{TE}}$, where p is the format precision, an invalid operation exception is recognized.

When the delivered result differs in value from the operand in FRB[p], an inexact exception is recognized. No underflow exception is recognized by this operation, regardless of the value of the operand in $\mathrm{FRB}[\mathrm{p}]$.
The FPSCR FPRF field is always set to the class and sign of the result, except for an enabled invalid-operation exception, in which case the field remains unchanged.

```
Special Registers Altered:
    FPRF FR FI
    FX XX
    VXSNAN VXCVI
    CR1
    (if Rc=1)
```


## Programming Note

DFP Quantize Immediate can be used to adjust values to a form having the specified exponent in the range -16 to 15 . If the adjustment requires the significand to be shifted left, then:

■ if the result would cause overflow from the most significant digit, the result is a default QNaN.;

- otherwise the result is the adjusted value (left shifted with matching exponent).
If the adjustment requires the significand to be shifted right, the result is rounded based on the value of the RMC field.

DFP Quantize Immediate can round a value to a specific number of fractional digits. Consider the computation of sales tax. Values expressed in U.S. dollars have 2 fractional digits, and sales tax rates typically have 3 fractional digits. The product of value and rate will yield 5 fractional digits. For example:

$$
39.95 * 0.075=2.99625
$$

This result needs to be rounded to the penny to compute the correct tax of $\$ 3.00$.

The following sequence computes the sales tax assuming the pre-tax total is in FRA and the tax rate is in FRB. The DFP Quantize Immediate instruction rounds the product (FRA * FRB) to 2 fractional digits (TE field $=-2$ ) using Round to nearest, ties away from 0 (RMC field = 2). The quantized and rounded result is placed in FRT.

```
dmul f0,FRA,FRB
dquai -2,FRT,f0,2
```


## DFP Quantize [Quad]

| dqua | FRT,FRA,FRB,RMC | $(R c=0)$ |
| :--- | :--- | :--- |
| dqua. | FRT,FRA,FRB,RMC | $(R c=1)$ |


| 59 | FRT | FRA | FRB | RMC |  | 3 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 | 23 |  |
| 0 |  |  |  |  |  |  |

dquaq
dquaq.
FRTp,FRAp,FRBp,RMC (Rc=0)

| 63 | FRTp | FRAp | FRBp | RMC |  | 3 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 | 21 | 23 |  |

The DFP operand in register $\operatorname{FRB}[\mathrm{p}]$ is converted and rounded to the form with the same exponent as that of the DFP operand in FRA[p] based on the rounding mode specified in the RMC field. The result of that form is placed in FRT[p]. The sign of the result is the same as the sign of the operand in FRB[p]. The ideal exponent is the exponent specified in FRA[p].
When the value of the operand in $\operatorname{FRB}[\mathrm{p}]$ is greater than $\left(10^{\mathrm{p}}-1\right) \times 10^{\mathrm{Ea}}$, where p is the format precision and Ea is the exponent of the operand in FRA[p], an invalid operation exception is recognized.

When the delivered result differs in value from the operand in $\operatorname{FRB}[p]$, an inexact exception is recognized. No
underflow exception is recognized by this operation, regardless of the value of the operand in FRB[p].
Figure 90 and Figure 91 summarize the actions. The tables do not include the setting of the FPSCR FPRF field. The FPSCR FPRF field is always set to the class and sign of the result, except for an enabled invalid-operation exception, in which case the field remains unchanged.

## Special Register Altered:

## FPRF FR FI

FX XX
VXSNAN VXCVI
CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

DFP Quantize can be used to adjust one DFP value (FRB[p]) to a form having the same exponent as a second DFP value (FRA[p]). If the adjustment requires the significand to be shifted left, then:

■ if the result would cause overflow from the most significant digit, the result is a default QNaN.;

- otherwise the result is the adjusted value (left shifted with matching exponent).

If the adjustment requires the significand to be shifted right, the result is rounded based on the value of the RMC field. Figure 89 shows examples of these adjustments.

| FRA | FRB | FRT when RMC=1 | FRT when RMC=2 |
| :---: | :---: | :---: | :---: |
| $1\left(1 \times 10^{0}\right)$ | $9 .\left(9 \times 10^{0}\right)$ | $9\left(9 \times 10^{0}\right)$ | $9\left(9 \times 10^{0}\right)$ |
| $1.00\left(100 \times 10^{-2}\right)$ | $9 .\left(9 \times 10^{0}\right)$ | $9.00\left(900 \times 10^{-2}\right)$ | $9.00\left(900 \times 10^{-2}\right)$ |
| $1\left(1 \times 10^{0}\right)$ | $49.1234\left(491234 \times 10^{-4}\right)$ | $49\left(49 \times 10^{0}\right)$ | $49\left(49 \times 10^{0}\right)$ |
| $1.00\left(100 \times 10^{-2}\right)$ | $49.1234\left(491234 \times 10^{-4}\right)$ | $49.12\left(4912 \times 10^{-2}\right)$ | $49.12\left(4912 \times 10^{-2}\right)$ |
| $1\left(1 \times 10^{0}\right)$ | $49.9876\left(499876 \times 10^{-4}\right)$ | $49\left(49 \times 10^{0}\right)$ | $50\left(50 \times 10^{0}\right)$ |
| $1.00\left(100 \times 10^{-2}\right)$ | $49.9876\left(499876 \times 10^{-4}\right)$ | $49.98\left(4998 \times 10^{-2}\right)$ | $49.99\left(4999 \times 10^{-2}\right)$ |
| $0.01\left(1 \times 10^{-2}\right)$ | $49.9876\left(499876 \times 10^{-4}\right)$ | $49.98\left(4998 \times 10^{-2}\right)$ | $49.99\left(4999 \times 10^{-2}\right)$ |
| $1\left(1 \times 10^{0}\right)$ | 9999999999999999 <br> $\left(9999999999999999 \times 10^{0}\right)$ | 9999999999999999 |  |
| $\left(9999999999999999 \times 10^{0}\right)$ | $\left(9999999999999999 \times 10^{0}\right)$ |  |  |
| $1.0\left(10 \times 10^{-1}\right)$ | 9999999999999999 |  |  |
| $\left(9999999999999999 \times 10^{0}\right)$ | $9 N a 9999999$ |  |  |

Figure 89. DFP Quantize examples

| Operand a in FRA[p] is | Actions for Quantize when operand $b$ in FRB[p] is |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0 | Fn | $\infty$ | QNaN | SNaN |
| 0 | * | * | $\mathrm{V}_{\text {XCVI: }}$ T(dNaN) | P (b) | $V_{\text {XSNAN }}$ : U(b) |
| Fn | * | * | $\mathrm{V}_{\text {XCVI: }}$ T(dNaN) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}: ~ U(\mathrm{~b})$ |
| - | $\mathrm{V}_{\text {XCVI }}: \mathrm{T}(\mathrm{dNaN})$ | $\mathrm{V}_{\mathrm{XCVI}}$ : $\mathrm{T}(\mathrm{dNaN})$ | T(dINF) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| QNaN | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{P}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{b})$ |
| SNaN | $V_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $V_{\text {XSNAN }}: ~ U(a)$ | $V_{\text {XSNAN }}: ~ U(a)$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{a})$ |
| Explanation: |  |  |  |  |  |
|  | See next table. |  |  |  |  |
| dINF | Default infinity |  |  |  |  |
| dNaN | Default quiet NaN |  |  |  |  |
| Fn | Finite nonzero numbers (includes both subnormal and normal numbers) |  |  |  |  |
| $\mathrm{P}(\mathrm{x})$ | The QNaN of operand $x$ is propagated and placed in FRT[p] |  |  |  |  |
| T(x) | The value x is placed in FRT[p] |  |  |  |  |
| $\mathrm{U}(\mathrm{x})$ | The SNaN of operand $x$ is converted to the corresponding QNaN and placed in FRT[p]. |  |  |  |  |
| $\mathrm{V}_{\mathrm{XCVI}}$ | The Invalid-Operation Exception (VXCVI) occurs. The result is produced only when the exception is disabled. (See Section 5.5.10.1 for actions) |  |  |  |  |
| $\mathrm{V}_{\text {XSNAN }}$ | The Invalid-Operation Exception (VXSNAN) occurs. The result is produced only when the exception is disabled. (See Section 5.5.10.1 for actions) |  |  |  |  |

Figure 90. Actions (part 1) Quantize


Figure 91. Actions (part2) Quantize

## DFP Reround [Quad] Z23-form

| drrnd drrnd. | FRT,FRA,FRB,RMC FRT,FRA,FRB,RMC |  |  |  |  | $\begin{array}{r} (\mathrm{Rc}=0) \\ (\mathrm{Rc}=1 \end{array}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ${ }_{0}^{59}$ | ${ }_{6}$ FRT | ${ }_{11} \mathrm{FRA}$ | FRB <br> 16 | RMC <br> 21 |  | 35 | Rc <br> 31 |
| drrndq drrndq. | FRTp,FRA,FRBp,RMC FRTp,FRA,FRBp,RMC |  |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| $63$ | FRTp <br> 6 | FRA <br> 11 | FRBp <br> 16 | RMC <br> 21 | 23 | 35 | Rc <br> 31 |

Let k be the contents of bits 58:63 of FRA that specifies the reference significance.

When the DFP operand in FRB[p] is a finite number, and if the reference significance is zero, or if the reference significance is nonzero and the number of significant digits of the source operand is less than or equal to the reference significance, then the value and the form of the source operand is placed in FRT[p]. If the reference significance is nonzero and the number of significant digits of the source operand is greater than the reference significance, then the source operand is converted and rounded to the number of significant digits specified in the reference significance based on the rounding mode specified in the RMC field. The result of the form with the specified number of significant digits is placed in FRT[p]. The sign of the result is the same as the sign of the operand in FRB[p].

For this instruction, the number of significant digits of the value 0 is considered to be zero. The ideal exponent is the greater value of the exponent of the operand in FRB[p] and the referenced exponent. The referenced exponent is the resultant exponent if the operand in FRB[p] would have been converted and rounded to the number of significant digits specified in the reference significance based on the rounding mode specified in the RMC field.

If the exponent of the rounded result of the form that has the specified number of significant digits would be greater than $X_{\text {max }}$, an invalid operation exception (VXCVI) occurs. When the invalid-operation exception occurs, and if the exception is disabled, a default QNaN is returned. When an invalid-operation exception occurs, no inexact exception is recognized.

In the absence of an invalid-operation exception, if the result differs in value from the operand in FRB[p], an inexact exception is recognized.

This operation causes neither an overflow nor an underflow exception.

Figure 93 summarizes the actions for Reround. The table does not include the setting of the FPSCR FPRF field. The FPSCR FPRF field is always set to the class and sign of the result, except for an enabled
invalid-operation exception, in which case the field remains unchanged.

## Special Registers Altered:

FPRF FR FI
FX XX
VXSNAN VXCVI
CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

DFP Reround can be used to adjust a DFP value (FRB[p]) to have no more than a specified number (FRA[p]58:63) of significant digits. The result (FRT[p]) is right-justified leaving the specified number of digits and rounded as specified by the RMC field. If rounding increases the number of significant digits, the result is adjusted again (the significand is shifted right 1 digit and the exponent is incremented by 1). Figure 92 has example results from DFP Reround for 1,2 , and 10 significant digits.

## Programming Note

DFP Reround is primarily used to round a DFP value to a specific number of digits before conversion to string format for printing or display. Another use for DFP Reround is to obtain the effective exponent of the most significant digit by specifying a reference significance of 1 . The exponent can be extracted and used to compute the number of significant digits or to left-justify a value.

For example, the following sequence computes the number of significant digits and returns it as an integer. FRB is the DFP value for which we want the number of significant digits; f13 contains the reference significance value 0x0000000000000001; and r 1 is the stack pointer, with free space for doublewords at offsets -8 and -16 . These doublewords are used to transfer the biased exponents from the FPRs to GPRs for integer computation. R3 contains the result of $E($ reround $(1, F R A))-E(F R A)+1$, where $E(x)$ represents the biased exponent of $x$.

```
dxex f0,FRB
stfd f0,-16(r1)
drrnd f1,f13,FRB,1 # reround 1 digit toward 0
dxex f1,f1
stfd f1,-8(r1)
lfd r11,-16(r1)
lfd r3,-8(r1)
subf r3,r11,r3
addi r3,r3,1
```

Given the value 412.34 the result is $E(4 \times 102)$ $E(41234 \times 10-2)+1=(398+2)-(398-2)+1=400-$ $396+1=5$. Additional code is required to detect and handle special values like Subnormal, Infinity, and NAN.

| FRA $_{58: 63}$ (binary) | FRB | FRT when RMC=1 | FRT when RMC=2 |
| :---: | :---: | :---: | :---: |
| 1 | $0.41234\left(41234 \times 10^{-5}\right)$ | $0.4\left(4 \times 10^{-1}\right)$ | $0.4\left(4 \times 10^{-1}\right)$ |
| 1 | $4.1234\left(41234 \times 10^{-4}\right)$ | $4\left(4 \times 10^{0}\right)$ | $4\left(4 \times 10^{0}\right)$ |
| 1 | $41.234\left(41234 \times 10^{-3}\right)$ | $4\left(4 \times 10^{1}\right)$ | $4\left(4 \times 10^{1}\right)$ |
| 1 | $412.34\left(41234 \times 10^{-2}\right)$ | $4\left(4 \times 10^{2}\right)$ | $4\left(4 \times 10^{2}\right)$ |
| 2 | $0.491234\left(491234 \times 10^{-6}\right)$ | $0.49\left(49 \times 10^{-2}\right)$ | $0.49\left(49 \times 10^{-2}\right)$ |
| 2 | $0.499876\left(499876 \times 10^{-6}\right)$ | $0.49\left(49 \times 10^{-2}\right)$ | $0.50\left(50 \times 10^{-2}\right)$ |
| 2 | $0.999876\left(999876 \times 10^{-6}\right)$ | $0.99\left(99 \times 10^{-2}\right)$ | $1.0\left(10 \times 10^{-1}\right)$ |
| 10 | $0.491234\left(491234 \times 10^{-6}\right)$ | $0.491234\left(491234 \times 10^{-6}\right)$ | $0.491234\left(491234 \times 10^{-6}\right)$ |
| 10 | $999.999\left(999999 \times 10^{-3}\right)$ | $999.999\left(999999 \times 10^{-3}\right)$ | $999.999\left(999999 \times 10^{-3}\right)$ |
| 10 | 9999999999999999 <br> $\left(9999999999999999 \times 10^{0}\right)$ | $9.999999999 \mathrm{E}+14$ <br> $\left(9999999999 \times 10^{5}\right)$ | $1.000000000 \mathrm{E}+15$ <br> $\left(1000000000 \times 10^{6}\right)$ |

Figure 92. DFP Reround examples

## Programming Note

DFP Reround combined with DFP Quantize can be used to left justify a value (as needed by the frexp function). FRB is the DFP value for which we want to left justify; f 13 contains the reference significance value $0 \times 0000000000000001$; and $r 1$ is the stack pointer, with free space for a doubleword at offset -8. This doubleword is used to transfer the biased exponents from the FPR to a GPR, for integer computation. The adjusted biased exponent (+ format precision - 1) is transferred back into an FPR so it can be inserted into the rerounded value. The adjusted rerounded value becomes the quantize reference value. The quantize instruction returns the left justified result in FRT.

```
drrnd f1,f13,FRB,1 \# reround 1 digit toward 0
dxex f0,f1
stfd f0,-8(r1)
lfd r11,-8(r1)
addi r11,r11,15 \# biased exp + precision - 1
lfd r11,-8(r1)
stfd \(f 0,-8(r 1)\)
diex f1,f0,f1 \# adjust exponent
dqua \(F R T, f 1, f 0,1\) \# quantize to adjusted
    exponent
```


## Version 3.0

|  | Actions for Reround when operand $b$ in FRB[p] is |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0* | Fn | $\infty$ | QNaN | SNaN |
| $\mathbf{k} \neq \mathbf{0}, \mathrm{k}<\mathrm{m}$ | - | $\begin{gathered} \mathrm{RR}(\mathrm{~b}) \text { or } \\ \mathrm{V}_{\mathrm{XCVI}}: \mathrm{T}(\mathrm{dNaN}) \end{gathered}$ | T(dINF) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| $\mathrm{k} \neq 0, \mathrm{k}=\mathrm{m}$ |  | W(b) | T(dINF) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : U(b) |
| $\begin{gathered} k \neq 0 \text { and } k>m, \\ \quad \text { or } k=0 \end{gathered}$ | W(b) | W(b) | T(dINF) | $\mathrm{P}(\mathrm{b})$ | $\mathrm{V}_{\text {XSNAN }}$ : $\mathrm{U}(\mathrm{b})$ |
| Explanation: |  |  |  |  |  |
|  | The number of significant digits of the value 0 is considered to be zero for this instruction. |  |  |  |  |
|  | Not applicable. |  |  |  |  |
| dINF | Default infinity. |  |  |  |  |
| Fn | Finite nonzero numbers (includes both subnormal and normal numbers). |  |  |  |  |
| k | Reference significance, which specifies the number of significant digits in the target operand. |  |  |  |  |
| m | Number of significant digits in the operand in FRB[p]. |  |  |  |  |
| $\mathrm{P}(\mathrm{x})$ | The QNaN of operand $x$ is propagated and placed in FRT[p]. |  |  |  |  |
| RR(x) | The value x is rounded to the form that has the specified number of significant digits. If $R R(x) \leq\left(10^{k}-1\right) \times 10^{X \max }$, then $R R(x)$ is returned; otherwise an invalid-operation exception is recognized. |  |  |  |  |
| T(x) | The value x is placed in FRT[p]. |  |  |  |  |
| $\mathrm{U}(\mathrm{x})$ | The SNaN of operand x is converted to the corresponding QNaN and placed in FRT[p]. |  |  |  |  |
| $\mathrm{V}_{\mathrm{XCVI}}$ | The Invalid-Operation Exception (VXCVI) occurs. The result is produced only when the exception is disabled. (See Section 5.5.10.1 for actions.) |  |  |  |  |
| $\mathrm{V}_{\text {XSNAN }}$ : | The Invalid-Operation Exception (VXSNAN) occurs. The result is produced only when the exception is disabled. See Section 5.5.10.1 for actions. |  |  |  |  |
| W(x) | The value and the form of $x$ is placed in FRT[p]. |  |  |  |  |

Figure 93. Actions: Reround
DFP Round To FP Integer With Inexact [Quad] Z23-form

| drintx | R,FRT,FRB,RMC | $(R \mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| drintx. | R,FRT,FRB,RMC | $(\mathrm{Rc}=1)$ |


| 59 | FRT | I/I | R | FRB | RMC |  | 99 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 15 | 16 | 21 | 23 |  |


| drintxq | R,FRTp,FRBp,RMC | $($ Rc=0 $)$ |
| :--- | :--- | :--- |
| drintxq. | R,FRTp,FRBp,RMC | $(R c=1)$ |


| 63 | FRTp | I// | R | FRBp | RMC |  | 99 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 15 | 16 | 21 | 23 |  | 31 |

The DFP operand in $\operatorname{FRB}[p]$ is rounded to a float-ing-point integer and placed into FRT[p]. The sign of the result is the same as the sign of the operand in $\mathrm{FRB}[\mathrm{p}]$. The ideal exponent is the larger value of zero and the exponent of the operand in $\operatorname{FRB}[p]$.

The rounding mode used is specified in the RMC field. When the RMC-encoding-selection (R) bit is zero, the RMC field contains the primary encoding; when the bit is one, the field contains the secondary encoding.

In addition to coercion of the converted value to fit the target format, the special rounding used by Round To FP Integer also coerces the target exponent to the ideal exponent.

When the operand in $\operatorname{FRB}[p]$ is a finite number and the exponent is less than zero, the operand is rounded to the result with an exponent of zero. When the exponent is greater than or equal to zero, the result is set to the numerical value and the form of the operand in FRB[p].

When the result differs in value from the operand in FRB[p], an inexact exception is recognized. No underflow exception is recognized by this operation, regardless of the value of the operand in FRB[p].

Figure 94 summarizes the actions for Round To FP Integer With Inexact. The table does not include the setting of the FPSCR ${ }_{\text {FPRF }}$ field. The FPSCR FPRF field is always set to the class and sign of the result, except for an enabled invalid-operation, in which case the field remains unchanged.

```
Special Registers Altered:
    FPRF FR FI
    FX XX
    VXSNAN
    CR1
        (if Rc=1)
```


## Programming Note

The DFP Round To FP Integer With Inexact and DFP Round To FP Integer With Inexact Quad instructions can be used to implement the decimal equivalent of the C99 rint function by specifying the primary RMC encoding for round according to FPSCR $_{\text {DRN }}(\mathrm{R}=0, \mathrm{RMC}=11)$. The specification for rint requires the inexact exception be raised if detected.

| Operand b in FRB is | Is $\boldsymbol{n}$ not precise ( $n \neq b$ ) | Inv.-Op. Exception Enabled | Inexact Exception Enabled | Is $\mathbf{n}$ Incremented ( $\|\mathrm{nl}\|>\|\mathrm{lb}\|$ ) | Actions* |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $-\infty$ | No ${ }^{1}$ | - | - | - | $\mathrm{T}(-\mathrm{dINF}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| F | No |  | - |  | $\mathrm{W}(\mathrm{n}), \mathrm{FI} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| F | Yes | - | No | No | $\mathrm{W}(\mathrm{n}), \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1$ |
| F | Yes | - | No | Yes | $\mathrm{W}(\mathrm{n}), \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1$ |
| F | Yes |  | Yes | No | $\mathrm{W}(\mathrm{n}), \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| F | Yes | - | Yes | Yes | $\mathrm{W}(\mathrm{n}), \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| $+\infty$ | No ${ }^{1}$ | - | - | - | $\mathrm{T}(+\mathrm{dINF}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| QNaN | No ${ }^{1}$ | - | - | - | $\mathrm{P}(\mathrm{b}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| SNaN | No ${ }^{1}$ | No | - | - | $\mathrm{U}(\mathrm{b}), \mathrm{FI} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{VXSNAN} \leftarrow 1$ |
| SNaN | No ${ }^{1}$ | Yes | - | - | VXSNAN $\leftarrow 1$, TV |
| Explanation: |  |  |  |  |  |
| * | Setting of XX and VXSNAN is part of the corresponding exception actions. Also, when an invalid-operation exception occurs, setting of FI and FR is part of the exception actions.(See the sections, "Inexact Exception" and "Invalid Operation Exception" for more details.) |  |  |  |  |
| - | The actions do not depend on this condition. |  |  |  |  |
| 1 | This condition is true by virtue of the state of some condition to the left of this column. |  |  |  |  |
| dINF | Default infinity. |  |  |  |  |
| F | All finite numbers, including zeros. |  |  |  |  |
| FI | Floating-Point-Fraction-Inexact status flag, FPSCR ${ }_{\text {Fl }}$. |  |  |  |  |
| FR | Floating-Point-Fraction-Rounded status flag, FPSCR $_{\text {FR }}$. |  |  |  |  |
| n | The value derived when the source operand, $b$, is rounded to an integer using the special rounding for Round To FP Integer. |  |  |  |  |
| $\mathrm{P}(\mathrm{x})$ | The QNaN of operand x is propagated and placed in FRT[p]. |  |  |  |  |
| T(x) | The value $x$ is placed in FRT[p]. |  |  |  |  |
| TV | The system floating-point enabled exception error handler is invoked for the invalid-operation exception if the FE0 and FE1 bits in the machine-state register are set to any mode other than the ignore-exception mode. |  |  |  |  |
| TX | The system floating-point enabled exception error handler is invoked for the inexact exception if the FE0 and FE1 bits in the machine-state register are set to any mode other than the ignore-exception mode. |  |  |  |  |
| $\mathrm{U}(\mathrm{x})$ | The SNaN of operand $x$ is converted to the corresponding QNaN and placed in FPT[p]. |  |  |  |  |
| W(x) | The value $x$ in the form of zero exponent or the source exponent is placed in FRT[p]. |  |  |  |  |
| XX | Floating-Point-Inexact-Exception status flag, FPSCR ${ }_{\text {XX }}$. |  |  |  |  |

Figure 94. Actions: Round to FP Integer With Inexact

## DFP Round To FP Integer Without Inexact [Quad] <br> Z23-form

| drintn | R,FRT,FRB,RMC | $(R \mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| drintn. | R,FRT,FRB,RMC | $(\mathrm{Rc}=1)$ |


| 59 | FRT | I// | R | FRB | RMC | 227 | Rc |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 15 | 16 | 21 | 23 |  |


| drintnq | R,FRTp,FRBp,RMC | $(R c=0)$ |
| :--- | :--- | :--- |
| drintnq. | R,FRTp,FRBp,RMC | $(R c=1)$ |


| 63 | FRTp | I/I | R | FRBp | RMC |  | 227 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 15 | 16 | 21 | 23 |  | 31 |

This operation is the same as the Round To FP Integer With Inexact operation, except that this operation does not recognize an inexact exception.

Figure 95 summarizes the actions for Round To FP Integer Without Inexact. The table does not include the setting of the FPSCR ${ }_{\text {FPRF }}$ field. The FPSCR FPRFF field is always set to the class and sign of the result, except for an enabled invalid-operation, in which case the field remains unchanged.

## Special Registers Altered:

FPRF FR (set to 0) Fl (set to 0)
FX
VXSNAN
CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

The DFP Round To FP Integer Without Inexact and DFP Round To FP Integer Without Inexact Quad instructions can be used to implement decimal equivalents of several C99 rounding functions by specifying the appropriate $R$ and RMC field values.

| Function | R | RMC |
| :--- | :--- | :--- |
| Ceil | 1 | Ob00 |
| Floor | 1 | $0 b 01$ |
| Nearbyint | 0 | $0 b 11$ |
| Round | 0 | $0 b 10$ |
| Trunc | 0 | Ob01 |

Note that nearbyint is similar to the rint function but without raising the inexact exception. Similarly ceil, floor, round, and trunc do not require the inexact exception.

| Operand $\mathbf{b}$ in FRB is | Inv.-Op. Exception Enabled | Actions* |
| :---: | :---: | :---: |
| $-\infty$ | - | $\mathrm{T}(-\mathrm{dINF}), \mathrm{FI} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| F |  | $\mathrm{W}(\mathrm{n}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| $+\infty$ | - | $\mathrm{T}(+\mathrm{dINF}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| QNaN |  | $\mathrm{P}(\mathrm{b}), \mathrm{FI} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| SNaN | No | $\mathrm{U}(\mathrm{b}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{VXSNAN} \leftarrow 1$ |
| SNaN | Yes | VXSNAN $\leftarrow 1, \mathrm{TV}$ |
| Explanation: |  |  |
| Sett ex O | Setting of VXSNAN is part of the corresponding exception actions. Also, when an invalid-operation exception occurs, setting of FI and FR bits is part of the exception actions. (See the sections, "Invalid Operation Exception" for more details.) |  |
| The | The actions do not depend on this condition. |  |
| dINF Defa | Default infinity. |  |
| F All fi | All finite numbers, including zeros. |  |
| FI Floa | Floating-Point-Fraction-Inexact status flag, FPSCR ${ }_{\text {Fl }}$. |  |
| FR Floa | Floating-Point-Fraction-Rounded status flag, FPSCR $_{\text {FR }}$. |  |
| n The | The value derived when the source operand, $b$, is rounded to an integer using the special rounding for Round-To-FP-Integer. |  |
| $\mathrm{P}(\mathrm{x}) \quad$ The | The QNaN of operand x is propagated and placed in FRT[p]. |  |
| T(x) The | The value x is placed in FRT[p]. |  |
| TV The | The system floating-point enabled exception error handler is invoked for the invalid-operation exception if the FE0 and FE1 bits in the machine-state register are set to any mode other than the ignore-exception mode. |  |
| $U(\mathrm{x}) \quad$ The | The SNaN of operand $x$ is converted to the corresponding QNaN and placed in FPT[p]. |  |
| $W(x) \quad$ The | The value $x$ in the form of zero exponent or the source exponent is placed in FRT[p]. |  |

Figure 95. Actions: Round to FP Integer Without Inexact

### 5.6.5 DFP Conversion Instructions

The DFP conversion instructions consist of data-format conversion instructions and data-type conversion instructions. They are all X-form instructions and employ the record bit (Rc).

### 5.6.5.1 DFP Data-Format Conversion Instructions

The data-format conversion instructions consist of Convert To DFP Long, Convert To DFP Extended, Round To DFP Short, and Round To DFP Long. Figure 96 summarizes the actions for these instructions.

## Programming Note

DFP does not provide operations on short operands, so they must be converted to long format, and then converted back to be stored. Preserving correct signaling NaN semantics requires that signaling NaNs be propagated from the source to the result without recognizing an exception during widening from short to long or narrowing from long to short. Because DFP does not provide equivalents to the FP Load Floating-Point Single and Store Floating-Point Single functions, the widening is performed by loading the DFP short value with a Load Floating as Integer Word Indexed followed by a DFP Convert to DFP Long, and narrowing is performed by a DFP Round to DFP Short followed by a Store Floating-Point as Integer Word Indexed. If the SNaN or infinity in DFP short format uses the preferred DPD encoding, then converting this operand to DFP long format and back to DFP short will result in the original bit pattern.

| Instruction | Actions when operand b in FRB[p] is |  |  |  |
| :--- | :---: | :---: | :---: | :---: |
|  | $\mathbf{F}$ | $\infty$ | $\mathbf{Q N a N}$ | SNaN |
| Convert To DFP Long | $\mathrm{T}(\mathrm{b})^{1}$ | $\mathrm{P}(\mathrm{b})^{2,4}$ | $\mathrm{P}(\mathrm{b})^{2,4}$ | $\mathrm{P}(\mathrm{b})^{3,4}$ |
| Convert To DFP Extended | $\mathrm{T}(\mathrm{b})^{1}$ | $\mathrm{~T}(\mathrm{dINF})$ | $\mathrm{P}(\mathrm{b})^{2,4}$ | $\mathrm{~V}_{\mathrm{XSNAN}}: \mathrm{U}(\mathrm{b})^{2,4}$ |
| Round To DFP Short | $\mathrm{R}(\mathrm{b})^{1}$ | $\mathrm{P}(\mathrm{b})^{2,5}$ | $\mathrm{P}(\mathrm{b})^{2,5}$ | $\mathrm{P}(\mathrm{b})^{3,5}$ |
| Round To DFP Long | $\mathrm{R}(\mathrm{b})^{1}$ | $\mathrm{~T}(\mathrm{dINF})$ | $\mathrm{P}(\mathrm{b})^{2,5}$ | $\mathrm{~V}_{\mathrm{XSNAN}}: \mathrm{U}(\mathrm{b})^{2,5}$ |

## Explanation:

1The ideal exponent is the exponent of the source operand.
2 Bits $5: \mathrm{N}-1$ of the N -bit combination field are set to zero.
3Bit 5 of the N -bit combination field is set to one. Bits $6: \mathrm{N}-1$ of the combination field are set to zero.
4The trailing significand field is padded on the left with zeros.
5Leftmost digits in the trailing significand field are removed.
dINFDefault infinity.
FAll finite numbers, including zeros.
$P(x)$ The special symbol in operand $x$ is propagated into $F R T[p]$.
$R(x)$ The value $x$ is rounded to the target-format precision; see Section 5.5.11
$T(x)$ The value $x$ is placed in FRT[p].
$U(x)$ The SNaN of operand x is converted to the corresponding QNaN.
$V_{\text {XSNAN }}$ The Invalid-Operation Exception (VXSNAN) occurs. The result is produced only when the exception is disabled. See Section 5.5.10.1 for actions.

Figure 96. Actions: Data-Format Conversion Instructions


The DFP short operand in bits 32:63 of FRB is converted to DFP long format and the converted result is placed into FRT. The sign of the result is the same as the sign of the source operand. The ideal exponent is the exponent of the source operand.

If the operand in FRB is an SNaN , it is converted to an SNaN in DFP long format and does not cause an invalid-operation exception.

## Special Registers Altered:

| FPRF FR (undefined) FI (undefined) |  |
| :--- | :--- | :--- |
| CR1 |  |
| (if $\mathrm{Rc}=1)$ |  |

## - Programming Note

Note that DFP short format is a storage-only format, Therefore, conversion of a short SNaN to long format will not cause an exception and the SNaN is preserved. Subsequent operation on that SNaN in long format will cause an exception.

## DFP Convert To DFP Extended X-form



The DFP long operand in the FRB is converted to DFP extended format and placed into FRTp. The sign of the result is the same as the sign of the operand in FRB. The ideal exponent is the exponent of the operand in FRB.

If the operand in FRB is an SNaN , an invalid-operation exception is recognized. If the exception is disabled, the SNaN is converted to the corresponding QNaN in DFP extended format.

## Special Registers Altered:

FPRF FR (set to 0) Fl (set to 0)
FX
VXSNAN
CR1
(if $\mathrm{Rc}=1$ )

| drsp | FRT,FRB | $(\mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| drsp. | FRT,FRB | $(\mathrm{Rc}=1)$ |


| 59 | FRT | I/I | FRB |  | 770 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |
| 31 |  |  |  |  |  |

The DFP long operand in FRB is converted and rounded to DFP short format. The DFP short value is extended on the left with zeros to form a 64-bit entity and placed into FRT. The sign of the result is the same as the sign of the source operand. The ideal exponent is the exponent of the source operand.

If the operand in FRB is an SNaN , it is converted to an SNaN in DFP short format and does not cause an invalid-operation exception.

Normally, the result is in the format and length of the target. However, when an overflow or underflow exception occurs and if the exception is enabled, the operation is completed by producing a wrapped rounded result in the same format and length as the source but rounded to the target-format precision.

## Special Registers Altered:

## FPRF FR FI

FX OX UX XX
CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

Note that DFP short format is a storage-only format, Therefore, conversion of a long SNaN to short format will not cause an exception. Converting a long format SNaN to short format is an implied move operation.

| DFP Round To DFP Long |  |  |  | X-form |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| drdpq drdpq. | FRTp,FRBp <br> FRTp,FRBp |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| 06 | ${ }_{6}$ FRTp | $11 /$ | FRBp <br> 16 | 21770 | Rc <br> 31 |

The DFP extended operand in FRBp is converted and rounded to DFP long format. The result concatenated with 640 s is placed in FRTp. The sign of the result is the same as the sign of the source operand. The ideal exponent is the exponent of the operand in FRBp.
If the operand in FRBp is an SNaN, an invalid-operation exception is recognized. If the exception is disabled, the SNaN is converted to the corresponding QNaN in DFP long format.
Normally, the result is in the format and length of the target. However, when an overflow or underflow exception occurs and if the exception is enabled, the operation is completed by producing a wrapped rounded result in the same format and length as the source but rounded to the target-format precision.

## Special Registers Altered:

FPRF FR FI
FX OX UX XX
VXSNAN
CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

Note that DFP Round to DFP Long, while producing a result in DFP long format, actually targets a register pair, writing 640 s in FRTp+1.

### 5.6.5.2 DFP Data-Type Conversion Instructions

The DFP data-type conversion instructions are used to convert data type between DFP and fixed.

The data-type conversion instructions consist of Convert From Fixed and Convert To Fixed.

## DFP Convert From Fixed

## X-form

| dcffix <br> dcffix. | FRT,FRB FRT,FRB |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $59$ | ${ }_{6} \text { FRT }$ | ${ }_{11} / I /$ | $\begin{aligned} & \hline \text { FRB } \\ & 16 \end{aligned}$ | 802 | Rc <br> 31 |

The 64-bit signed binary integer in FRB is converted and rounded to a DFP Long value and placed into FRT. The sign of the result is the same as the sign of the source operand. The ideal exponent is zero.

If the source operand is a zero, then a plus zero with a zero exponent is returned.

The FPSCR FPRF field is set to the class and sign of the result.

```
Special Registers Altered:
    FPRF FR FI
    FX XX
    CR1
(if \(\mathrm{Rc}=1\) )
```


## DFP Convert From Fixed Quad

X-form

| dcffixq | FRTp,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| dcffixq. | FRTp,FRB | $(R c=1)$ |


| 63 | FRTp | I/I | FRB |  | 802 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 |  |

The 64-bit signed binary integer in FRB is converted and rounded to a DFP Extended value and placed into FRTp. The sign of the result is the same as the sign of the source operand. The ideal exponent is zero.
If the source operand is a zero, then a plus zero with a zero exponent is returned.

The FPSCR FPRRF field is set to the class and sign of the result.

## Special Registers Altered:

$$
\begin{aligned}
& \text { FPRF FR (undefined) } \quad \mathrm{Fl} \text { (undefined) } \\
& \text { CR1 }
\end{aligned}
$$

## DFP Convert To Fixed [Quad] X-form

| detfix | FRT,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| dctfix. | FRT,FRB | $(R c=1)$ |


| 59 | FRT | I/I | FRB |  | 290 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |


| dctfixq | FRT,FRBp | $(R c=0)$ |
| :--- | :--- | :--- |
| dctfixq. | FRT,FRBp | $(R c=1)$ |


| 63 | FRT | I/I | FRBp |  | 290 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | Rc |  |  |

The DFP operand in FRB[p] is rounded to an integer value and is placed into FRT in the 64-bit signed binary integer format. The sign of the result is the same as the sign of the source operand, except when the source operand is a NaN or a zero.

Figure 97 summarizes the actions for Convert To Fixed.

```
Special Registers Altered:
    FPRF (undefined) FR FI
    FX XX
    VXSNAN VXCVI
    CR1
```

                                    (if \(\mathrm{Rc}=1\) )
    
## Programming Note

It is recommended that software pre-round the operand to a floating-point integral using drintx[q] or drintn[ $q$ ] is a rounding mode other than the current rounding mode specified by FPSCR $_{\text {DRN }}$ is needed. Saving, modifying and restoring the FPSCR just to temporarily change the rounding mode is less efficient than just employing drintx[p] or drint[ $p$ ] which override the current rounding mode using an immediate control field.

For example if the desired function rounding is Round to Nearest, Ties away from 0 but the default rounding (from FPSCR ${ }_{\text {DRN }}$ ) is Round to Nearest, Ties to Even then following is preferred.

```
drintn 0,£1,f1,2
dctfix f1,f1
```

| Operand b in FRB[p] is | $q$ is | Is n not precise ( $\mathbf{n} \neq \mathbf{b}$ ) | Inv.-Op. Except. Enabled | Inexact Except. Enabled | Is $\mathbf{n}$ Incremented ( $\|n\|>\|b\|)$ | Actions * |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $-\infty \leq \mathrm{b}<\mathrm{MN}$ | < MN | - | No | - | - | $\mathrm{T}(\mathrm{MN}), \mathrm{FI} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{VXCVI} \leftarrow 1$ |
| $-\infty \leq \mathrm{b}<\mathrm{MN}$ | < MN | - | Yes | - | - | $\mathrm{VXCVI} \leftarrow 1$, TV |
| $-\infty<\mathrm{b}<\mathrm{MN}$ | = MN | - | - | No | - | $\mathrm{T}(\mathrm{MN}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1$ |
| $-\infty<\mathrm{b}<\mathrm{MN}$ | $=\mathrm{MN}$ | - | - | Yes | - | $\mathrm{T}(\mathrm{MN}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| $\mathrm{MN} \leq \mathrm{b}<0$ | - | No | - | - | - | $\mathrm{T}(\mathrm{n}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| $\mathrm{MN} \leq \mathrm{b}<0$ | - | Yes | - | No | No | $\mathrm{T}(\mathrm{n}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1$ |
| $\mathrm{MN} \leq \mathrm{b}<0$ | - | Yes | - | No | Yes | $\mathrm{T}(\mathrm{n}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1$ |
| $\mathrm{MN} \leq \mathrm{b}<0$ | - | Yes | - | Yes | No | $\mathrm{T}(\mathrm{n}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| $\mathrm{MN} \leq \mathrm{b}<0$ | - | Yes | - | Yes | Yes | $\mathrm{T}(\mathrm{n}), \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| $\pm 0$ | - | No | - | - | - | $\mathrm{T}(0), \mathrm{FI} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| $0<\mathrm{b} \leq \mathrm{MP}$ | - | No | - | - | - | $\mathrm{T}(\mathrm{n}), \mathrm{FI} \leftarrow 0, \mathrm{FR} \leftarrow 0$ |
| $0<\mathrm{b} \leq \mathrm{MP}$ | - | Yes | - | No | No | $\mathrm{T}(\mathrm{n}), \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1$ |
| $0<\mathrm{b} \leq \mathrm{MP}$ | - | Yes | - | No | Yes | $\mathrm{T}(\mathrm{n}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1$ |
| $0<\mathrm{b} \leq \mathrm{MP}$ | - | Yes | - | Yes | No | $\mathrm{T}(\mathrm{n}), \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| $0<\mathrm{b} \leq \mathrm{MP}$ | - | Yes | - | Yes | Yes | $\mathrm{T}(\mathrm{n}), \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 1, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| MP < b < + | = MP | - | - | No | - | $\mathrm{T}(\mathrm{MP}), \mathrm{Fl} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1$ |
| MP < b < + | = MP | - | - | Yes | - | $\mathrm{T}(\mathrm{MP}), \mathrm{FI} \leftarrow 1, \mathrm{FR} \leftarrow 0, \mathrm{XX} \leftarrow 1, \mathrm{TX}$ |
| MP $<$ b $\leq+\infty$ | > MP | - | No | - | - | $\mathrm{T}(\mathrm{MP}), \mathrm{FI} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{VXCVI} \leftarrow 1$ |
| $\mathrm{MP}<\mathrm{b} \leq+\infty$ | > MP | - | Yes | - | - | $\mathrm{VXCVI} \leftarrow 1$, TV |
| QNaN | - |  | No | - | - | $\mathrm{T}(\mathrm{MN}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{VXCVI} \leftarrow 1$ |
| QNaN | - | - | Yes | - | - | VXCVI $\leftarrow 1$, TV |
| SNaN | - | - | No | - | - | $\mathrm{T}(\mathrm{MN}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{VXCVI} \leftarrow 1, \mathrm{VXSNAN} \leftarrow 1$ |
| SNaN | - | - | Yes | - | - | $\mathrm{VXCVI} \leftarrow 1, \mathrm{VXSNAN} \leftarrow 1$, TV |
| Explanation: |  |  |  |  |  |  |
|  | Setting of XX, VXCVI, and VXSNAN is part of the corresponding exception actions. Also, when an invalid-operation exception occurs, setting of FI and FR bits is part of the exception actions. (See the sections, "Inexact Exception" and "Invalid Operation Exception" for more details.) |  |  |  |  |  |
| T | The actions do not depend on this condition. |  |  |  |  |  |
| FI F | Floating-Point-Fraction-Inexact status flag, FPSCR $_{\text {FI }}$. |  |  |  |  |  |
| FR F | Floating-Point-Fraction-Rounded status flag, FPSCR $_{\text {FR }}$. |  |  |  |  |  |
| MN M | Maximum negative number representable by the 64-bit binary integer format |  |  |  |  |  |
| MP M | Maximum positive number representable by the 64-bit binary integer format. |  |  |  |  |  |
| n T | The value q converted to a fixed-point result. |  |  |  |  |  |
| $\mathrm{q} \quad \mathrm{T}$ | The value derived when the source value $b$ is rounded to an integer using the specified rounding mode |  |  |  |  |  |
| $\mathrm{T}(\mathrm{x}) \quad \mathrm{T}$ | The value x is placed in FRT[p]. |  |  |  |  |  |
| TV T | The system floating-point enabled exception error handler is invoked for the invalid-operation exception if the FE0 and FE1 bits in the machine-state register are set to any mode other than the ignore-exception mode. |  |  |  |  |  |
| $\text { TX } \quad \text { T }$ | The system floating-point enabled exception error handler is invoked for the inexact exception if the FE0 and FE1 bits in the machine-state register are set to any mode other than the ignore-exception mode. |  |  |  |  |  |
| VXCVI T | The FPSCR ${ }_{\text {VXCVI }}$ invalid operation exception status bit. |  |  |  |  |  |
| VXSNAN T | The FPSCR ${ }_{\text {VXSNAN }}$ invalid operation exception status bit. |  |  |  |  |  |
| XX F | Floating-Point-Inexact-Exception status flag, FPSCR $_{\text {XX }}$. |  |  |  |  |  |

Figure 97. Actions: Convert To Fixed

### 5.6.6 DFP Format Instructions

The DFP format instructions are used to compose or decompose a DFP operand. A source operand of SNaN does not cause an invalid-operation exception. All format instructions employ the record bit (Rc).

The format instructions consist of Decode DPD To $B C D$, Encode BCD To DPD, Extract Biased Exponent, Insert Biased Exponent, Shift Significand Left Immediate, and Shift Significand Right Immediate.

## DFP Decode DPD To BCD [Quad] X-form

| ddedpd | SP,FRT,FRB | $(R \mathrm{c}=0)$ |
| :--- | :--- | :--- |
| ddedpd. | SP,FRT,FRB | $(R \mathrm{c}=1)$ |


| 59 | FRT | SP <br> 0 | I/I | FRB |  | 322 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 11 | 13 | 21 |  | 31 |  |  |  |

ddedpdq ddedpdq

| 63 | FRTp | SP | I/I | FRBp |  | 322 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 13 | 16 | 21 |  |  |

A portion of the significand of the DFP operand in $\mathrm{FRB}[\mathrm{p}]$ is converted to a signed or unsigned BCD number depending on the SP field. For infinity and NaN , the significand is considered to be the contents in the trailing significand field padded on the left by a zero digit.

## SP $\mathbf{O}_{0}=\mathbf{0}$ (unsigned conversion)

The rightmost 16 digits of the significand (32 digits for ddedpdq) is converted to an unsigned BCD number and the result is placed into FRT[p].

## $S P_{0}=1$ (signed conversion)

The rightmost 15 digits of the significand ( 31 digits for ddedpdq) is converted to a signed BCD number with the same sign as the DFP operand, and the result is placed into FRT[p]. If the DFP operand is negative, the sign is encoded as 0b1101. If the DFP operand is positive, $\mathrm{SP}_{1}$ indicates which preferred plus sign encoding is used. If $S P_{1}=0$, the plus sign is encoded as $0 b 1100$ (the option-1 preferred sign code), otherwise the plus sign is encoded as Ob1111(the option-2 preferred sign code).

## Special Registers Altered:

CR1
(if $\mathrm{Rc}=1$ )

## DFP Encode BCD To DPD [Quad] X-form

| denbcd | S,FRT,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| denbcd. | S,FRT,FRB | $(R c=1)$ |


| 59 FRT S I/I FRB  834  <br> 0 6 11 12 16 21   <br> 31        |
| :--- |


| 63 | FRTp | S | I/I | FRBp |  | 834 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 12 | 16 | 21 |  | 31 |

The signed or unsigned BCD operand, depending on the $S$ field, in FRB[p] is converted to a DFP number. The ideal exponent is zero.

## S = 0 (unsigned BCD operand)

The unsigned BCD operand in FRB[p] is converted to a positive DFP number of the same magnitude and the result is placed into FRT[p].

## S = 1 (signed BCD operand)

The signed BCD operand in $\operatorname{FRB}[p]$ is converted to the corresponding DFP number and the result is placed into FRT[p].

If an invalid BCD digit or sign code is detected in the source operand, an invalid-operation exception (VXCVI) occurs.

FPSCR $_{\text {FPRF }}$ is set to the class and sign of the result, except for Invalid Operation Exception when FPSCR ${ }_{V E}=1$.

## Special Registers Altered:

FPRF FR (set to 0) FI (set to 0)
FX
VXCVI
CR1

## DFP Extract Biased Exponent [Quad] <br> X-form

| dxex | FRT,FRB | $(R c=0)$ |
| :--- | :--- | :--- |
| dxex. | FRT,FRB | $(R c=1)$ |


| 59 | FRT |  | FRB |  | 354 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |


| dxexq | FRT,FRBp | $(R c=0)$ |
| :--- | :--- | :--- |
| dxexq. | FRT,FRBp | $(R c=1)$ |


| 63 | FRT | I/I | FRBp |  | 354 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |

The biased exponent of the operand in $\mathrm{FRB}[\mathrm{p}]$ is extracted and placed into FRT in the 64-bit signed binary integer format. When the operand in FRB is an infinity, QNaN, or SNaN , a special code is returned.

| Operand | Result |
| :--- | :--- |
| Finite Number | biased exponent value |
| Infinity | -1 |
| QNaN | -2 |
| SNaN | -3 |

## Special Registers Altered:

CR1
(if $\mathrm{Rc}=1$ )

## Programming Note

The exponent bias value is 101 for DFP Short, 398 for DFP Long, and 6176 for DFP Extended.

## DFP Insert Biased Exponent [Quad] X-form

| diex | FRT,FRA,FRB | $(R \mathrm{Rc}=0)$ |
| :--- | :--- | :--- |
| diex. | FRT,FRA,FRB | $(\mathrm{Rc}=1)$ |


| 59 | FRT | FRA <br> 0 | FRB <br> 16 | 21 | 866 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 61 |  |  |  |  |  |

$$
\begin{array}{lll}
\text { diexq } & \text { FRTp,FRA,FRBp } & (\mathrm{Rc}=0) \\
\text { diexq. } & \text { FRTp,FRA,FRBp } & (R \mathrm{Rc}=1)
\end{array}
$$

| 63 | FRTp | FRA | FRBp |  | 866 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 11 | 16 | 21 |  |

Let $a$ be the value of the 64-bit signed binary integer in FRA.
a $\quad$ Result
$a>M B E^{1} \quad$ QNaN
$M B E \geq a \geq 0$ Finite number with biased exponent a
$a=-1 \quad$ Infinity
$a=-2 \quad$ QNaN
$a=-3 \quad S N a N$
$a<-3 \quad$ QNaN
1 Maximum biased exponent for the target format
When $0 \leq a \leq$ MBE, $a$ is the biased target exponent that is combined with the sign bit and the significand value of the DFP operand in FRB[p] to form the DFP result in FRT[p]. The ideal exponent is the specified target exponent.
When a specifies a special code ( $\mathrm{a}<0$ or $\mathrm{a}>\mathrm{MBE}$ ), an infinity, QNaN, or SNaN is formed in FRT[p] with the trailing significand field containing the value from the trailing significand field of the source operand in FRB[p], and with an N -bit combination field set as follows.

- For an Infinity result,

■ the leftmost 5 bits are set to 0b11110, and

- the rightmost $\mathrm{N}-5$ bits are set to zero.
- For a QNaN result,

■ the leftmost 5 bits are set to 0b11111,

- bit 5 is set to zero, and
- the rightmost $\mathrm{N}-5$ bits are set to zero.
- For an SNaN result,

■ the leftmost 5 bits are set to Ob11111,

- bit 5 is set to one, and
- the rightmost $\mathrm{N}-5$ bits are set to zero.


## Special Registers Altered:

CR1
(if $R c=1$ )

## Programming Note

The exponent bias value is 101 for DFP Short, 398 for DFP Long, and 6176 for DFP Extended.

| Operand a in FRA[p] specifies | Actions for Insert Biased Exponent when operand b in FRB[p] specifies |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
|  | F | $\infty$ | QNaN | SNaN |
| F | N, Rb | Z, Rb | Z, Rb | Z, Rb |
| $\infty$ | I, Rb | I, Rb | I, Rb | I, Rb |
| QNaN | Q, Rb | Q, Rb | Q, Rb | Q, Rb |
| SNaN | S, Rb | S, Rb | S, Rb | S, Rb |
| Explanation: |  |  |  |  |
| F | All finite numbers, including zeros |  |  |  |
| 1 | The combination field in FRT[p] is set to indicate a default Infinity. |  |  |  |
| N | The combination field in FRT[p] is set to the specified biased exponent in FRA and the leftmost significand digit in FRB[p]. |  |  |  |
| Q | The combination field in FRT[p] is set to indicate a default QNaN. |  |  |  |
| S | The combination field in FRT[p] is set to indicate a default SNaN. |  |  |  |
| Z | The combination field in FRT[p] is set to indicate the specific biased exponent in FRA and a leftmost coefficient digit of zero. |  |  |  |
| Rb | The contents of the trailing significand field in FRB[p] are reencoded using preferred DPD encodings and the reencoded result is placed in the same field in FRT[p]. The sign bit of FRB[p] is copied into the sign bit in FRT[p]. |  |  |  |

Figure 98. Actions: Insert Biased Exponent

## DFP Shift Significand Left Immediate [Quad] <br> Z22-form

| dscli dscli. | $\begin{aligned} & \text { FRT,FRA,SH } \\ & \text { FRT,FRA,SH } \end{aligned}$ |  |  | $\begin{aligned} & (R c=0) \\ & (R c=1) \end{aligned}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $59$ | FRT 6 | ${ }_{11} \mathrm{FRA}$ | ${ }_{16} \mathrm{SH}$ | 2266 | Rc <br> 31 |
| dscliq dscliq. | FRTp,FRAp,SH <br> FRTp,FRAp,SH |  |  |  | $\begin{aligned} & (R c=0) \\ & (R c=1) \end{aligned}$ |
| 063 | FRTp <br> 6 | FRAp | ${ }_{16} \mathrm{SH}$ | 2266 | Rc <br> 31 |

The significand of the DFP operand in FRA[p] is shifted left SH digits. For a NaN or infinity, all significand digits are in the trailing significand field. SH is a 6-bit unsigned binary integer. Digits shifted out of the leftmost digit are lost. Zeros are supplied to the vacated positions on the right. The result is placed into FRT[p]. The sign of the result is the same as the sign of the source operand in FRA[p].

If the source operand in FRA[p] is a finite number, the exponent of the result is the same as the exponent of the source operand.
For an Infinity, QNaN or SNaN result, the target format's N -bit combination field is set as follows.

- For an Infinity result,

■ the leftmost 5 bits are set to 0b11110, and

- the rightmost $\mathrm{N}-5$ bits are set to zero.
- For a QNaN result,
- the leftmost 5 bits are set to 0b11111,
- bit 5 is set to zero, and
- the rightmost $\mathrm{N}-6$ bits are set to zero.
- For an SNaN result,

■ the leftmost 5 bits are set to 0b11111,

- bit 5 is set to one, and
- the rightmost $\mathrm{N}-6$ bits are set to zero.


## Special Registers Altered:

## CR1

(if $R c=1$ )

DFP Shift Significand Right Immediate [Quad] Z22-form

| dscri dscri. | FRT,FRA,SH <br> FRT,FRA,SH |  |  |  |  | $\begin{aligned} & (\mathrm{Rc}=0) \\ & (\mathrm{Rc}=1) \end{aligned}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $59$ | ${ }_{6}$ FRT | $\begin{aligned} & \hline \text { FRA } \\ & 11 \end{aligned}$ | ${ }_{16} \mathrm{SH}$ | 22 | 98 | $\left\|\begin{array}{l}\mathrm{Rc} \\ 31\end{array}\right\|$ |
| dscriq dscriq. | FRT | ,FRAp, |  | FRTp,FRAp,SH |  | $\begin{aligned} & (R c=0) \\ & (R c=1) \end{aligned}$ |


| 63 | FRTp | FRAp | SH |  | 98 | Rc |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 | 11 | 16 | 22 |  | 31 |

The significand of the DFP operand in FRA[p] is shifted right SH digits. For a NaN or infinity, all significand digits are in the trailing significand field. SH is a 6 -bit unsigned binary integer. Digits shifted out of the units digit are lost. Zeros are supplied to the vacated positions on the left. The result is placed into FRT[p]. The sign of the result is the same as the sign of the source operand in FRA[p].
If the source operand in FRA[p] is a finite number, the exponent of the result is the same as the exponent of the source operand.
For an Infinity, QNaN or SNaN result, the target format's N -bit combination field is set as follows.

- For an Infinity result, ■ the leftmost 5 bits are set to Ob11110, and
- the rightmost $\mathrm{N}-5$ bits are set to zero.
- For a QNaN result,

■ the leftmost 5 bits are set to 0b11111,

- bit 5 is set to zero, and
- the rightmost $\mathrm{N}-6$ bits are set to zero.
- For an SNaN result,

■ the leftmost 5 bits are set to Ob11111,

- bit 5 is set to one, and
- the rightmost $\mathrm{N}-6$ bits are set to zero.


## Special Registers Altered:

CR1
(if $\mathrm{Rc}=1$ )

### 5.6.7 DFP Instruction Summary

|  | Full Name | $\begin{aligned} & \sum_{\substack{r}}^{0} \\ & \text { O} \end{aligned}$ | Operands | SNaN <br> Vs G |  | FPRF |  | $\begin{gathered} \text { FP } \\ \text { Exception } \\ \text { V Z O U X } \end{gathered}$ |  | IE | U |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  | $\cup$ | U U H |  |  |  |  |
| dadd | DFP Add | X | FRT, FRA, FRB | Y N | RE | Y | Y | V O U X | Y | Y | Y |
| daddq | DFP Add Quad | X | FRTp, FRAp, FRBp | Y N | RE | Y | Y | $V \quad O \cup X$ | Y | Y | Y |
| dsub | DFP Subtract | X | FRT, FRA, FRB | Y N | RE | Y | Y | $V \quad O \cup X$ | Y | Y | Y |
| dsubq | DFP Subtract Quad | X | FRTp, FRAp, FRBp | Y N | RE | Y | Y | $V \quad O \cup X$ | Y | Y | Y |
| dmul | DFP Multiply | X | FRT, FRA, FRB | Y N | RE | Y | Y | $V \quad O \cup X$ | Y | Y | Y |
| dmulq | DFP Multiply Quad | X | FRTp, FRAp, FRBp | Y N | RE | Y | Y | $V \quad O \cup X$ | Y | Y | Y |
| ddiv | DFP Divide | X | FRT, FRA, FRB | Y N | RE | Y | Y | V Z O U X | Y | Y | Y |
| ddivq | DFP Divide Quad | X | FRTp, FRAp, FRBp | Y N | RE | Y | Y | V Z O U X | Y | Y | Y |
| dcmpo | DFP Compare Ordered | X | BF, FRA, FRB | Y - | - | N | Y | V | - | - | N |
| dcmpoq | DFP Compare Ordered Quad | X | BF, FRAp, FRBp | Y | - | N | Y | V | - | - | N |
| dcmpu | DFP Compare Unordered | X | BF, FRA, FRB | Y | - | N | Y | V | - | - | N |
| dcmpuq | DFP Compare Unordered Quad | X | BF, FRAp, FRBp | Y | - | N | Y | V | - | - | N |
| dtstdc | DFP Test Data Class | Z22 | BF, FRA, DCM | N - | - | N | $Y^{1}$ |  | - | - | N |
| dtstdcq | DFP Test Data Class Quad | Z22 | BF, FRAp, DCM | N | - | N | $Y^{1}$ |  | - | - | N |
| dtstdg | DFP Test Data Group | Z22 | BF, FRA,DGM | N | - | N | $Y^{1}$ |  | - | - | N |
| dtstdgq | DFP Test Data Group Quad | Z22 | BF, FRAp, DGM | N | - | N | $Y^{1}$ |  | - | - | N |
| dtstex | DFP Test Exponent | X | BF, FRA, FRB | N | - | N | Y |  | - | - | N |
| dtstexq | DFP Test Exponent Quad | X | BF, FRAp, FRBp | N | - | N | Y |  | - | - | N |
| dtstsf | DFP Test Significance | X | BF, FRA(FIX), FRB | N | - | N | Y |  | - | - | N |
| dtstsfq | DFP Test Significance Quad | X | BF, FRA(FIX), FRBp | N | - | N | Y |  | - | - | N |
| dquai | DFP Quantize Immediate | Z23 | TE, FRT, FRB, RMC | Y N | RE | Y | Y | $V \quad \mathrm{X}$ | Y | Y | Y |
| dquaiq | DFP Quantize Immediate Quad | Z23 | TE, FRTp, FRBp, RMC | Y N | RE | Y | Y | $V \quad \mathrm{X}$ | Y | Y | Y |
| dqua | DFP Quantize | Z23 | FRT,FRA,FRB,RMC | Y N | RE | Y | Y | $V \quad \mathrm{X}$ | Y | Y | $Y$ |
| dquaq | DFP Quantize Quad | Z23 | FRTp,FRAp,FRBp, RMC | Y N | RE | Y | Y | V X | Y | Y | $Y$ |
| drrnd | DFP Reround | Z23 | FRT,FRA(FIX),FRB,RMC | Y N | RE | Y | Y | V X | Y | Y | $Y$ |
| drrndq | DFP Reround Quad | Z23 | FRTp, FRA(FIX), FRBp, RMC | Y N | RE | Y | Y | V X | Y | Y | $Y$ |
| drintx | DFP Round To FP Integer With Inexact | Z23 | R,FRT, FRB,RMC | Y N | RE | Y | Y | $V \quad \mathrm{X}$ | Y | Y | Y |
| drintxq | DFP Round To FP Integer With Inexact Quad | Z23 | R,FRTp,FRBp,RMC | Y N | RE | Y | Y | V X | Y | Y | Y |
| drintn | DFP Round To FP Integer Without Inexact | Z23 | R,FRT, FRB,RMC | Y N | RE | Y | Y | V | $\mathrm{Y}^{\#}$ | Y | Y |
| drintnq | DFP Round To FP Integer Without Inexact Quad | Z23 | R,FRTp, FRBp,RMC | Y N | RE | Y | Y | V | $\mathrm{Y}^{\#}$ | Y | Y |
| dctdp | DFP Convert To DFP Long | X | FRT, FRB (DFP Short) | N Y | RE | Y | $Y^{2}$ |  | U | Y | Y |
| dctqpq | DFP Convert To DFP Extended | X | FRTp, FRB | Y N | RE | Y | Y | V | $\mathrm{Y}^{\#}$ | Y | Y |
| drsp | DFP Round To DFP Short | X | FRT (DFP Short), FRB | N Y | RE | Y | $Y^{2}$ | O UX | Y | Y | $Y$ |
| drdpq | DFP Round To DFP Long | X | FRTp, FRBp | Y N | RE | Y | Y | V O U X | Y | Y | $Y$ |
| dcffixq | DFP Convert From Fixed Quad | X | FRTp, FRB (FIX) | - N | RE | Y | Y |  | U | Y | Y |
| detfix | DFP Convert To Fixed | X | FRT (FIX), FRB | Y N | - | U | U | $V \quad \mathrm{X}$ | Y | - | $Y$ |
| dctfixq | DFP Convert To Fixed Quad | X | FRT (FIX), FRBp | Y N | - | U | U | $V \quad \mathrm{X}$ | Y | - | $Y$ |
| ddedpd | DFP Decode DPD To BCD | X | SP, FRT(BCD), FRB | N - | - | N | N |  | - | - | $Y$ |

Figure 99. Decimal Floating-Point Instructions Summary

|  | Full Name | $\begin{aligned} & \sum_{\text {ru}}^{0} \\ & \text { O} \end{aligned}$ | Operands | SNaN <br> Vs G | $\begin{aligned} & \text { 을 } \\ & \text { 응 } \\ & \text { U } \\ & \hline \end{aligned}$ | FPRF |  | $\begin{gathered} \text { FP } \\ \text { Exception } \\ \text { V Z O U X } \end{gathered}$ | $\frac{\overline{\text { İ }}}{\text { 号 }}$ | IE | U |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  | 0 | U O H |  |  |  |  |
| ddedpdq | DFP Decode DPD To BCD Quad | X | SP, FRTp(BCD), FRBp | N - | - | N | N |  | - | - | Y |
| denbcd | DFP Encode BCD To DPD | X | S, FRT, FRB (BCD) | - N | RE | Y | Y | V | $\mathrm{Y}^{\#}$ | Y | Y |
| denbcdq | DFP Encode BCD To DPD Quad | X | S, FRTp, FRBp (BCD) | - N | RE | Y | Y | V | $\mathrm{Y}^{\#}$ | Y | Y |
| dxex | DFP Extract Biased Exponent | X | FRT (FIX), FRB | N N | - | N | N |  | - | - | Y |
| dxexq | DFP Extract Biased Exponent Quad | X | FRT (FIX), FRBp | N N | - | N | N |  | - | - | Y |
| diex | DFP Insert Biased Exponent | X | FRT, FRA(FIX), FRB | N Y | RE | N | N |  | - | Y | Y |
| diexq | DFP Insert Biased Exponent Quad | X | FRTp, FRA(FIX), FRBp | N Y | RE | N | N |  | - | Y | Y |
| dscli | DFP Shift Significand Left Immediate | Z22 | FRT,FRA,SH | N Y | RE | N | N |  | - | - | Y |
| dscliq | DFP Shift Significand Left Immediate Quad | Z22 | FRTp,FRAp,SH | N Y | RE | N | N |  | - | - | Y |
| dscri | DFP Shift Significand Right Immediate | Z22 | FRT,FRA,SH | N Y | RE | N | N |  | - | - | Y |
| dscriq | DFP Shift Significand Right Immediate Quad | Z22 | FRTp,FRAp,SH | N Y | RE | N | N |  | - | - | Y |


| Explanation: |  |
| :---: | :---: |
| \# | FI and FR are set to zeros for these instructions. |
| - | Not applicable. |
| 1 | A unique definition of the FPSCR ${ }_{\text {FPCC }}$ field is provided for the instruction. |
| 2 | These are the only instructions that may generate an SNaN and also set the FPSC FPRR field. Since the BFP FPSCR FPRF field does not include a code for SNaN, these instructions cause the need for redefining the FPSCR FPRF field for DFP. |
| DCM | A 6-bit immediate operand specifying the data-class mask. |
| DGM | A 6-bit immediate operand specifying the data-group mask. |
| G | An SNaN can be generated as the target operand. |
| IE | An ideal exponent is defined for the instruction. |
| FI | Setting of the FPSCR ${ }_{\text {FI }}$ flag. |
| FR | Setting of the FPSCR ${ }_{\text {FR }}$ flag. |
| N | No. |
| O | An overflow exception may be recognized. |
| Rc | The record bit, Rc, is provided to record FPSCR ${ }_{32: 35}$ in CR field 1. |
| RE | The trailing significand field is reencoded using preferred DPD encodings. The preferred DPD encoding are also used for propagated NaNs , or converted NaNs and infinities. |
| RMC | A 2-bit immediate operand specifying the rounding-mode control. |
| S | An one-bit immediate operand specifying if the operation is signed or unsigned. |
| SP | A two-bit immediate operand: one bit specifies if the operation is signed or unsigned and, for signed operations, another bit specifies which preferred plus sign code is generated. |
| U | An underflow exception may be recognized. |
| V | An invalid-operation exception may be recognized. |
| Vs | An input operand of SNaN causes an invalid-operation exception. |
| X | An inexact exception may be recognized. |
| Y | Yes. |
| U | Undefined |
| Z | A zero-divide exception may be recognized. |

Figure 99. Decimal Floating-Point Instructions Summary (Continued)

## Chapter 6. Vector Facility

### 6.1 Vector Facility Overview

This chapter describes the registers and instructions that make up the Vector Facility.

### 6.2 Chapter Conventions

### 6.2.1 Description of Instruction

 OperationThe following notation, in addition to that described in Section 1.3.2, is used in this chapter. Additional RTL functions are described in Appendix C.

## x.bit[y]

Return the contents of bit $y$ of $x$.
x.bit[y:z]

Return the contents of bits $y: z$ of $x$.
x.nibble[y]

Return the contents of the 4-bit nibble element $y$ of $x$.
x.nibble[y:z]

Return the contents of the nibble elements $y: z$ of x.
x.byte[y]

Return the contents of byte element $y$ of $x$.
x.byte[y:z]

Return the contents of byte elements $y: z$ of $x$.
x.hword[y]

Return the contents of halfword element $y$ of $x$.

## x.hword[y:z]

Return the contents of halfword elements $y: z$ of $x$.

## x.word[y]

Return the contents of word element $y$ of $x$.
x.word[y:z]

Return the contents of word element $\mathrm{y}: \mathrm{z}$ of x .
x.dword[y]

Return the contents of doubleword element $y$ of $x$.
x.dword[y:z]

Return the contents of doubleword elements $y: z$ of x.
x?y:z
if the value of $x$ is true, then the value of $y$, otherwise the value $z$.
${ }^{+}$int
Integer addition.
${ }^{+}{ }_{f}$
Floating-point addition.
-fp
Floating-point subtraction.
${ }^{x_{\text {sui }}}$
Multiplication of a signed-integer (first operand) by an unsigned-integer (second operand).
$x_{f p}$
Floating-point multiplication.
$=_{\text {int }}$
Integer equals relation.
$=_{f p}$
Floating-point equals relation.
$<_{u i}, \leq_{u i},>_{u i}, \geq_{u i}$
Unsigned-integer comparison relations.
$<_{s i}, \leq_{s i},>_{s i}, \geq_{s i}$
Signed-integer comparison relations.
$<_{f p}, \leq_{f p},>_{f p}, \geq_{f p}$
Floating-point comparison relations.

## LENGTH( x )

Length of $x$, in bits. If $x$ is the word "element", $\operatorname{LENGTH}(x)$ is the length, in bits, of the element implied by the instruction mnemonic.
$x \ll y$
Result of shifting $x$ left by $y$ bits, filling vacated bits with zeros.

```
b \leftarrowLENGTH(x)
result \leftarrow(y<b) ? (xy:b-1 |>0): b}
```

$x \gg{ }_{\text {ui }} y$
Result of shifting $x$ right by $y$ bits, filling vacated bits with zeros.

```
b}\leftarrow\operatorname{LENGTH(x)
```


$x \gg y$
Result of shifting $x$ right by $y$ bits, filling vacated bits with copies of bit 0 (sign bit) of $x$.

$$
\begin{aligned}
& b \leftarrow \operatorname{LENGTH}(x) \\
& \text { result } \leftarrow(y<b) ?\left({ }^{y} x_{0} \| x_{0:(b-y) \cdot 1}\right):{ }^{b} x_{0}
\end{aligned}
$$

## $x \lll<y$

Result of rotating $x$ left by $y$ bits.

$$
\begin{aligned}
& b \leftarrow \operatorname{LENGTH}(x) \\
& \text { result } \leftarrow x_{y: b} \cdot 1 \| x_{0: y-1}
\end{aligned}
$$

## $x \ggg y$

Returns the contents of $x$ rotated right by $y$ bits.

## Chop( $\mathrm{x}, \mathrm{y}$ )

Result of extending the right-most $y$ bits of $x$ on the left with zeros.

```
result \leftarrowx& ((1<<y)\cdot1)
```


## Clamp( $x, y, z$ )

$x$ is interpreted as a signed integer. If the value of $x$ is less than $y$, then the value $y$ is returned, else if the value of $x$ is greater than $z$, the value $z$ is returned, else the value x is returned.

```
if (x<y) then
    result \leftarrowy
    VSCR SAT }\leftarrow
else if (x>z) then
    result }\leftarrow
    VSCR SAT }\leftarrow
else result }\leftarrow
```


## ConvertSItoBCD(x,y)

Let x be a signed integer quadword.
Let $y$ indicate the preferred sign code.
Return the signed integer value $x$ in packed decimal format.

```
if (x<0) then do
    x}\leftarrow\negx+
    sign}\leftarrow0\times000
end
else
    sign}\leftarrow(y=0) ? 0x000C: 0x000
result }\leftarrow
shont }\leftarrow
do while (x > 0)
    digit }\leftarrowx%1
    result }\leftarrow\mathrm{ result | (digit<<shcnt)
    x}\leftarrowx\div1
    shent & shent + 4
end
return(result | sign)
```


## ConvertBCDtoSI(x)

Let x be a packed decimal value.
Return the value $x$ in signed integer format.

```
result }\leftarrow
scale}\leftarrow
sign }\leftarrowx.bit[124:127]
x <x >> 4
do while (x>0)
    digit }\leftarrowx&0\times000
    result \leftarrow result + (digit x scale)
    x \leftarrowx >> 4
    scale}\leftarrow\mathrm{ scale x }1
end
if (sign==0\times000B) | (sign==0x000D) then
    result \leftarrow result + 1
return result
```


## ConvertSPtoSXWsaturate( $x, y$ )

Let x be a single-precision floating-point value.
Let $y$ be an unsigned integer value.

```
sign }\leftarrowx.bit[0
exp }\leftarrowx.\mathrm{ bit[1:8]
frac,bit[0:22]}\leftarrowx.bit[9:31]
frac.bit[23:30]}\leftarrow0b00000000
if (exp==255) & (frac!=0) then return (0x0000_0000) || NaN operand
if (exp==255) & (frac==0) then do | Infinity operand
VSCR.SAT \leftarrow1
        return ((sign==1) ? Ox8000_0000: Ox7FFF_FFFF)
end
if ((exp+Y-127)>30) then do | |arge operand
        VSCR.SAT \leftarrow1
        return ((sign==1) ? Ox8000_0000: Ox7FFF_FFFF)
end
if ((exp+y-127)<0) then return (0x0000_0000) | | . 1.0 < value < 1.0 (value rounds to 0)
significand.bit[0] \leftarrowObl
significand.bit[1:31] &frac
do i = 1 to 31.(exp+Y-127)
        significand \leftarrow significand >> uil
end
return ((sign==0) ? significand: ( ssignificand + l))
```


## ConvertSPtoUXWsaturate( $x, y$ )

Let x be a single-precision floating-point value.
Let $y$ be an unsigned integer value.

```
sign }\leftarrowx.bit[0
exp }\leftarrowx.bit[1:8
frac,bit[0:22]}\leftarrowx.bit[9:31]
frac,bit[23:30]}\leftarrow0b0000_000
if (exp==255) & (frac!=0) then return (0x0000_0000) |/ NaN operand
if (exp==255) & (frac==0) then do | Infinity operand
    VSCR.SAT \leftarrow1
    return((sign==1) ? Ox0000_0000: OxFFFF_FFFF)
end
if ((exp+Y-127)>31) then do || |arge operand
    VSCR.SAT \leftarrow }
    return((sign==1) ? Ox0000_0000: OxFFFF_FFFF)
end
if ((exp+Y-127)<0) then return (0x0000_0000)
if( sign==1 ) then do
1| - 1.0 < value < 1.0
|/ value rounds to 0
|/ negative operand
    VSCR.SAT \leftarrow1
    return (0x0000_0000)
end
significand.bit[0]}\leftarrow\mathrm{ obl
significand,bit[1:31]}\leftarrowfra
do i = 1 to 31.(exp+Y-127)
    significand = significand >>ui 1
end
return(significand)
```


## Version 3.0

## ConvertSXWtoSP(x)

Let x be a 32-bit signed integer value.

```
sign }\leftarrow\mathrm{ X.bit[0]
exp }\leftarrow32+12
frac.bit[0] \leftarrowx.bit[0]
frac.bit[1:32]}\leftarrowx.bit[0:31
if (frac==0) return (0x0000_0000) || Zero Operand
if (sign==1) then frac = ffrac + 1
do while (frac.bit[0]=0)
            frac \leftarrowfrac<< l
            exp & exp - 1
end
lsb \leftarrowfrac.bit[23]
gbit \leftarrowfrac.bit[24]
xbit \leftarrowfrac.bit[25:32]!=0
inc \leftarrow(lsb & gbit)| (gbit & xbit)
frac.bit[0:23]}\leftarrowfrac.bit[0:23] + in
if (carry_out=1) then exp \leftarrow exp + 1
result.bit[0]}\leftarrow\mathrm{ sign
result.bit[1:8]}\leftarrowex
result.bit[9:31] \leftarrowfrac.bit[1:23]
return (result)
```


## ConvertUXWtoSP(x)

```
exp \leftarrow < 31 + 127
frac}\leftarrowx.bit[0:31
if (frac==0) return (0x0000_0000) || Zero Operand
do while( fraco==0 )
    frac \leftarrowfrac<< l
    exp \leftarrowexp - 1
end
lsb \leftarrowfrac.bit[23]
gbit \leftarrowfrac.bit[24]
xbit \leftarrowfrac,bit[25:31]!=0
inc}\leftarrow(|sb& gbit)|(gbit & xbit
frac.bit[0:23]}\leftarrowfrac.bit[0:23] + in
if (carry_out=1) then exp \leftarrow exp + 1
result.bit[0] \leftarrowObo
result.bit[1:8]}\leftarrow\mathrm{ exp
result.bit[9:31] \leftarrowfrac.bit[1:23]
return (result)
```


## DUP( $x, y$ )

Return the concatenation of $y$ copies $x$.

```
DUP(ObO1,4) = Ob01010101
DUP(ObOO1,3) = 0b001001001
```


## EXTZ(x)

Result of extending $x$ on the left with zeros.

```
b}\leftarrow\operatorname{LENGTH(x)
result \leftarrowx& ((1<<b).1)
```


## InvMixColumns(x)

do $\mathrm{c}=0$ to 3
result.word[c].byte[0] = Ox0E•x.word[c].byte[0] ^ OxOB•x.word[c].byte[1] ^ OxOD•x.word[c].byte[2] ^ 0x09•x.word[c].byte[3] result.word[c].byte[1] $=0 \times 09 \bullet x$. word[c].byte[0] ^ $0 x 0 E \cdot x \cdot w o r d[c] . b y t e[1] ~ \wedge ~ O x 0 B \bullet x . w o r d[c] . b y t e[2] ~ \wedge ~ 0 x 0 D \cdot x . w o r d[c] . b y t e[3] ~$ result.word[c].byte[2] = 0x0D•x.word[c].byte[0] ^ 0x09•x.word[c].byte[1] ^ 0x0E•x.word[c].byte[2] ^ 0x0B•x.word[c].byte[3] result.word[c].byte[3] = Ox0B•x.word[c].byte[0] ^ Ox0D•x.word[c].byte[1] ^ 0x09•x.word[c].byte[2] ^ Ox0E•x.word[c].byte[3] end
return(result);
where "•" is a $\operatorname{GF}\left(2^{8}\right)$ multiply, a binary polynomial multiplication reduced by modulo $0 \times 11 \mathrm{~B}$.
The $G F\left(2^{8}\right)$ multiply of $0 \times 09 \cdot x$ can be expressed in minimized terms as the following. product.bit[0] $=x . b i t[0] \wedge x . b i t[3]$
product.bit[1] = x.bit[1] ^x.bit[4] ^x.bit[0]
product.bit[2] $=x . \operatorname{bit}[2] \wedge x . b i t[5] \wedge x . b i t[0] \wedge x . b i t[1]$
product.bit[3] = x.bit[3] ^x.bit[6] ^ x.bit[1] ^x.bit[2]
product.bit[4] = x.bit[4] ^ x.bit[7] ^ x.bit[0] ^x.bit[2]
product.bit[5] = x.bit[5] ^x.bit[0] ^x.bit[1]
product.bit[6] $=x . b i t[6] \wedge x . b i t[1] \wedge x . b i t[2]$
product.bit[7] = x.bit[7] ^ x.bit[2]
The $G F\left(2^{8}\right)$ multiply of $O \times O B \cdot \times$ can be expressed in minimized terms as the following.

```
product.bit[0] = x.bit[0] ^ x.bit[1] ^ x.bit[3]
product.bit[1] = x.bit[1] ^ x.bit[2] ^ x.bit[4] ^ x.bit[0]
product.bit[2] = x.bit[2] ^ x.bit[3] ^ x.bit[5] ^ x.bit[0] ^ x.bit[1]
product.bit[3] = x.bit[3] ^ x.bit[4] ^ x.bit[6] ^ x.bit[0] ^ x.bit[1] ^ x.bit[2]
product.bit[4] = x.bit[4] ^ x.bit[5] ^ x.bit[7] ^ x.bit[2]
product.bit[5] = x.bit[5] ^ x.bit[6] ^ x.bit[0] ^ x.bit[1]
product.bit[6] = x.bit[6] ^ x.bit[7] ^ x.bit[0] ^ x.bit[1] ^ x.bit[2]
product.bit[7] = x.bit[7] ^ x.bit[0] ^ x.bit[2]
```

The $G F\left(2^{8}\right)$ multiply of $0 \times O D \cdot \times$ can be expressed in minimized terms as the following.

```
product.bit[0] = x.bit[0] ^ x.bit[2] ^ x.bit[3]
product.bit[1] = x.bit[1] ^ x.bit[3] ^ x.bit[4] ^ x.bit[0]
product.bit[2] = x.bit[2] ^ x.bit[4] ^ x.bit[5] ^ x.bit[1]
product.bit[3] = x.bit[3] ^ x.bit[5] ^ x.bit[6] ^ x.bit[0] ^ x.bit[2]
product.bit[4] = x.bit[4] ^ x.bit[6] ^ x.bit[7] ^ x.bit[0] ^ x.bit[1] ^ x.bit[2]
product.bit[5] = x.bit[5] ^ x.bit[7] ^ x.bit[1]
product.bit[6] = x.bit[6] ^ x.bit[0] ^ x.bit[2]
product.bit[7] = x.bit[7] ^ x.bit[1] ^ x.bit[2]
```

The $G F\left(2^{8}\right)$ multiply of $0 \times 0 E \cdot \times$ can be expressed in minimized terms as the following.
product.bit[0] = x.bit[1] ^x.bit[2] ^x.bit[3]
product.bit[1] = x.bit[2] ^ x.bit[3] ^ x.bit[4] ^ x.bit[0]
product.bit[2] = x.bit[3] ^ x.bit[4] ^ x.bit[5] ^ x.bit[1]
product.bit[3] $=x . b i t[4] \wedge x . b i t[5] \wedge x . b i t[6] \wedge x . b i t[2]$
product.bit[4] = x.bit[5] ^ x.bit[6] ^ x.bit[7] ^ x.bit[1] ^ x.bit[2]
product.bit[5] = x.bit[6] ^ x.bit[7] ^ x.bit[1]
product.bit[6] = x.bit[7] ^ x.bit[2]
product.bit[7] = x.bit[0] ^ x.bit[1] ^ x.bit[2]

## Version 3.0

## InvShiftRows(x)

```
result.word[0].byte[0] = x.word[0].byte[0]
result.word[1].byte[0] = x.word[1].byte[0]
result.word[2].byte[0] = x.word[2].byte[0]
result.word[3].byte[0] = x.word[3].byte[0]
result.word[0].byte[1] = x.word[3].byte[1]
result.word[1].byte[1] = x.word[0].byte[1]
result.word[2].byte[1] = x.word[1].byte[1]
result.word[3].byte[1] = x.word[2].byte[1]
result.word[0].byte[2] = x.word[2].byte[2]
result.word[1].byte[2] = x.word[3].byte[2]
result.word[2].byte[2] = x.word[0].byte[2]
result.word[3].byte[2] = x.word[1].byte[2]
result.word[0].byte[3] = x.word[1].byte[3]
result.word[1].byte[3] = x.word[2].byte[3]
result.word[2].byte[3] = x.word[3].byte[3]
result.word[3].byte[3] = x.word[0].byte[3]
return(result)
```


## InvSubBytes(x)

InvSB0X.byte[256] $=\{0 x 52,0 x 09,0 x 6 A, 0 x D 5,0 x 30,0 \times 36,0 x A 5,0 x 38,0 x B F, 0 x 40,0 x A 3,0 x 9 E, 0 x 81,0 x F 3,0 x D 7,0 x F B$, $0 x 7 C, 0 x E 3,0 x 39,0 x 82,0 x 9 B, 0 x 2 F, 0 x F F, 0 x 87,0 x 34,0 x 8 E, 0 x 43,0 x 44,0 x C 4,0 x D E, 0 x E 9,0 x C B$, $0 x 54,0 \times 7 \mathrm{~B}, 0 \times 94,0 \times 32,0 \times A 6,0 x C 2,0 \times 23,0 \times 3 \mathrm{D}, 0 \mathrm{xEE}, 0 \mathrm{x} 4 \mathrm{C}, 0 \mathrm{x} 95,0 \mathrm{x} 0 \mathrm{~B}, 0 \mathrm{x} 42,0 \mathrm{xFA}, 0 \mathrm{xC} 3,0 \mathrm{x} 4 \mathrm{E}$, $0 x 08,0 x 2 E, 0 x A 1,0 x 66,0 x 28,0 x D 9,0 \times 24,0 x B 2,0 x 76,0 x 5 B, 0 x A 2,0 x 49,0 x 6 D, 0 x 8 B, 0 x D 1,0 x 25$, $0 x 72,0 \times F 8,0 \times F 6,0 \times 64,0 x 86,0 \times 68,0 \times 98,0 \times 16,0 \times D 4,0 \times 44,0 \times 5 C, 0 \times C C, 0 \times 5 D, 0 x 65,0 \times B 6,0 \times 92$, $0 x 6 C, 0 x 70,0 x 48,0 \times 50,0 x F D, 0 x E D, 0 x B 9,0 x D A, 0 x 5 E, 0 x 15,0 x 46,0 x 57,0 x A 7,0 x 8 D, 0 x 9 D, 0 x 84$, $0 x 90,0 x D 8,0 x A B, 0 x 00,0 x 8 C, 0 x B C, 0 x D 3,0 x 0 A, 0 x F 7,0 x E 4,0 x 58,0 x 05,0 x B 8,0 x B 3,0 x 45,0 x 06$, $0 x D 0,0 x 2 C, 0 x 1 E, 0 \times 8 F, 0 x C A, 0 x 3 F, 0 x 0 F, 0 x 02,0 x C 1,0 x A F, 0 x B D, 0 x 03,0 x 01,0 x 13,0 x 8 \mathrm{~A}, 0 \times 6 \mathrm{~B}$, $0 x 3 A, 0 x 91,0 x 11,0 x 41,0 x 4 F, 0 x 67,0 x D C, 0 x E A, 0 x 97,0 x F 2,0 x C F, 0 x C E, 0 x F 0,0 x B 4,0 x E 6,0 x 73$, $0 \times 96,0 \times A C, 0 \times 74,0 \times 22,0 \times E 7,0 \times A D, 0 \times 35,0 \times 85,0 \times E 2,0 \times F 9,0 \times 37,0 \times E 8,0 \times 1 C, 0 \times 75,0 \times D F, 0 \times 6 E$, $0 \times 47,0 \times F 1,0 x 1 A, 0 \times 71,0 \times 1 D, 0 \times 29,0 x C 5,0 \times 89,0 \times 6 F, 0 x B 7,0 x 62,0 \times 0 E, 0 \times A A, 0 x 18,0 x B E, 0 x 1 B$, $0 \times F C, 0 x 56,0 x 3 E, 0 x 4 B, 0 x C 6,0 x D 2,0 x 79,0 x 20,0 x 9 A, 0 x D B, 0 x C 0,0 x F E, 0 x 78,0 x C D, 0 x 5 A, 0 x F 4$, $0 \times 1 \mathrm{~F}, 0 \times \mathrm{DD}, 0 \times \mathrm{AA} 8,0 \times 33,0 \times 88,0 \times 07,0 \times C 7,0 \times 31,0 \times B 1,0 \times 12,0 \times 10,0 \times 59,0 \times 27,0 \times 80,0 \times E C, 0 \times 5 \mathrm{~F}$, $0 x 60,0 x 51,0 x 7 F, 0 x A 9,0 x 19,0 x B 5,0 x 4 A, 0 x 0 D, 0 x 2 D, 0 x E 5,0 x 7 \mathrm{~A}, 0 \mathrm{x} 9 \mathrm{~F}, 0 \mathrm{x} 93,0 \mathrm{xC} 9,0 \mathrm{x} 9 \mathrm{C}, 0 \mathrm{xEF}$, $0 \times A 0,0 \times E 0,0 \times 3 B, 0 x 4 D, 0 x A E, 0 \times 2 A, 0 \times F 5,0 \times B 0,0 x C 8,0 x E B, 0 x B B, 0 x 3 C, 0 \times 83,0 x 53,0 x 99,0 \times 61$, $0 \times 17,0 \times 2 B, 0 \times 04,0 \times 7 E, 0 \times B A, 0 \times 77,0 \times D 6,0 \times 26,0 \times E 1,0 \times 69,0 \times 14,0 \times 63,0 \times 55,0 \times 21,0 \times 0 C, 0 \times 7 D\}$

```
do i = 0 to 15
    result.byte[i] = InvSBOX.byte[x.byte[i]]
end
return(result)
```


## MixColumns(x)

do $\mathrm{c}=0$ to 3
result.word[c].byte[0] $=0 x 02 \cdot x \cdot w o r d[c] . b y t e[0] \wedge 0 x 03 \cdot x \cdot w o r d[c] . b y t e[1] \wedge \quad x \cdot w o r d[c] . b y t e[2] \wedge \quad x . w o r d[c] . b y t e[3]$ result.word[c].byte[1] $=\quad x . w o r d[c] . b y t e[0] ~ \wedge ~ 0 x 02 \cdot x . w o r d[c] . b y t e[1] ~ \wedge ~ 0 x 03 \cdot x . w o r d[c] . b y t e[2] ~ \wedge ~ x . w o r d[c] . b y t e[3] ~$ result.word[c].byte[2] $=\quad x$. word[c].byte[0] ^ $\quad x . w o r d[c] . b y t e[1] \wedge 0 x 02 \cdot x . w o r d[c] . b y t e[2] \wedge 0 x 03 \cdot x . w o r d[c] . b y t e[3]$ result.word[c].byte[3] $=0 x 03 \cdot x . w o r d[c] . b y t e[0] \wedge \quad x . w o r d[c] . b y t e[1] ~ \wedge \quad x . w o r d[c] . b y t e[2] ~ \wedge ~ 0 x 02 \cdot x . w o r d[c] . b y t e[3]$
end
return(result)
The $G F\left(2^{8}\right)$ multiply of $0 \times 02 \cdot \times$ can be expressed in minimized terms as the following.

```
    product.bit[0] = x.bit[1]
    product.bit[1] = x.bit[2]
    product.bit[2] = x.bit[3]
    product.bit[3] = x.bit[4] ^ x.bit[0]
    product.bit[4] = x.bit[5] ^ x.bit[0]
    product.bit[5] = x.bit[6]
    product.bit[6] = x.bit[7] ^ x.bit[0]
    product.bit[7] = x.bit[0]
```

The $G F\left(2^{8}\right)$ multiply of $0 \times 03 \cdot \times$ can be expressed in minimized terms as the following.

```
product.bit[0] = x.bit[0] ^ x.bit[1]
product.bit[1] = x.bit[1] ^ x.bit[2]
product.bit[2] = x.bit[2] ^ x.bit[3]
product.bit[3] = x.bit[3] ^ x.bit[4] ^ x.bit[0]
product.bit[4] = x.bit[4] ^ x.bit[5] ^ x.bit[0]
product.bit[5] = x.bit[5] ^ x.bit[6]
product.bit[6] = x.bit[6] ^ x.bit[7] ^ x.bit[0]
product.bit[7] = x.bit[7] ^ x.bit[0]
```


## ShiftRows(x)

```
result.word[0].byte[0] = x.word[0].byte[0]
result.word[1].byte[0] = x.word[1].byte[0]
result.word[2].byte[0] = x.word[2].byte[0]
result.word[3].byte[0] = x.word[3].byte[0]
result.word[0].byte[1] = x.word[1].byte[1]
result.word[1].byte[1] = x.word[2].byte[1]
result.word[2].byte[1] = x.word[3].byte[1]
result.word[3].byte[1] = x.word[0].byte[1]
result.word[0].byte[2] = x.word[2].byte[2]
result.word[1].byte[2] = x.word[3].byte[2]
result.word[2].byte[2] = x.word[0].byte[2]
result.word[3].byte[2] = x.word[1].byte[2]
result.word[0].byte[3] = x.word[3].byte[3]
result.word[1].byte[3] = x.word[0].byte[3]
result.word[2].byte[3] = x.word[1].byte[3]
result.word[3].byte[3] = x.word[2].byte[3]
return(result)
```


## Signed_BCD_Add(x,y,z)

Let $x$ and $y$ be 31-digit signed decimal values.
Performs a signed decimal addition of $x$ and $y$.
If the unbounded result is equal to zero, $e q_{-} f l a g$ is set to 1 . Otherwise, $e q_{-} f l a g$ is set to 0 . If the unbounded result is greater than zero, gt $f \mathrm{fl}$ ag is set to 1 . Otherwise, gt fl ag is set to 0 . If the unbounded result is less than zero, $\mid t_{-} f^{-} a_{g}$ is set to 1 . Otherwise, $\left|t_{-} f\right|_{\text {a }}^{-} g$ is set to 0 .

If the magnitude of the unbounded result is greater than $10^{31} \cdot 1,0 x_{-} f l a g$ is set to 1 . Otherwise, $0 x_{-} f l a g$ is set to 0.

If the unbounded result is greater than or equal to zero, the sign code of the result is set to $0 b 1100$ if $z=0$. If the unbounded result is greater than or equal to zero, the sign code of the result is set to 0 b 1111 if $z=1$. If the unbounded result is less than zero, the sign code of the result is set to 0 b1101.

The low-order 31 digits of the unbounded result magnitude concatented with the sign code are returned.
If either operand is an invalid encoding of a signed decimal value, the result returned is undefined and $\mathrm{inv} f \mid \operatorname{lag}$ is set to 1 and $\left|t_{\_} f\right| a g, g t_{\_} f \mid a g$ andeq_flag are set to 0 . Otherwise, $i_{n v} f \mid a g$ is set to 0 .

## Signed_BCD_Subtract(x,y,z)

Let $x$ and $y$ be 31-digit signed decimal values.
Performs a signed decimal subtract of $y$ from $x$.
If the unbounded result is equal to zero, $\mathrm{eq} q_{-} f \mathrm{ag}$ is set to 1 . Otherwise, eq $f l a g$ is set to 0 . If the unbounded result is greater than zero, gt flag is set to 1 . Otherwise, $g \mathrm{ft}_{\mathrm{f}} \mathrm{fl}$ ag is set to 0 .
If the unbounded result is less than zero, $\mid t_{-} f^{-} \mathrm{a}_{\mathrm{g}}$ is set to 1 . Otherwise, $\mathrm{It}_{-} \mathrm{fl}^{-} \mathrm{a}_{\mathrm{g}}$ is set to 0 .
If the magnitude of the unbounded result is greater than $10^{31} .1,0 x_{-} f \mathrm{lag}$ is set to 1 . Otherwise, $0 \mathrm{x}_{-} \mathrm{fl} \mathrm{ag}$ is set to 0 .

If the unbounded result is greater than or equal to zero, the sign code of the result is set to $0 b 1100$ if $z=0$. If the unbounded result is greater than or equal to zero, the sign code of the result is set to $0 b 1111$ if $z=1$. If the unbounded result is less than zero, the sign code of the result is set to 0 b1101.

The low-order 31 digits of the unbounded result magnitude concatented with the sign code are returned.
If either operand is an invalid encoding of a signed decimal value, the result returned is undefined and inv_flag


## SubBytes(x)



```
do i = 0 to 15
    result.byte[i] = SBOX.byte[x.byte[i]]
end
return(result)
```


## RoundToSPIntCeil(x)

The value $x$ if $x$ is a single-precision floating-point integer; otherwise the smallest single-precision floating-point integer that is greater than $x$.

## RoundToSPIntFloor(x)

The value x if x is a single-precision floating-point integer; otherwise the largest single-precision floating-point integer that is less than $x$.

## RoundToSPIntNear(x)

The value x if x is a single-precision floating-point integer; otherwise the single-precision floating-point integer that is nearest in value to $x$ (in case of a tie, the even single-precision floating-point integer is used).

## RoundToSPIntTrunc(x)

The value $x$ if $x$ is a single-precision floating-point integer; otherwise the largest single-precision floating-point integer that is less than $x$ if $x>0$, or the smallest single-precision floating-point integer that is greater than x if $\mathrm{x}<0$.

## RoundToNearSP(x)

The single-precision floating-point number that is nearest in value to the infinitely-precise floating-point intermediate result $x$ (in case of a tie, the single-precision floating-point value with the least-significant bit equal to 0 is used).

ReciprocalEstimateSP(x)
A single-precision floating-point estimate of the reciprocal of the single-precision floating-point number $x$.

ReciprocalSquareRootEstimateSP(x)
A single-precision floating-point estimate of the reciprocal of the square root of the single-precision floating-point number $x$.

## LogBase2EstimateSP(x)

A single-precision floating-point estimate of the base 2 logarithm of the single-precision floating-point number $x$.

## Power2EstimateSP(x)

A single-precision floating-point estimate of the 2 raised to the power of the single-precision floating-point number $x$.

### 6.3 Vector Facility Registers

I

| . qword |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| . dword[0] |  |  |  |  |  |  |  | . dword[ 1] |  |  |  |  |  |  |  |  |
| word [0] |  |  |  | . word [ 1] |  |  |  | . word [ 2] |  |  |  |  | word [ 3] |  |  |  |
| . hword[0] |  | hword [ 1] |  | hword [ 2] |  | . hword[3] |  | hword [ 4] |  | hword [ 5] |  |  | hword[ 6] |  | hword [ 7] |  |
| . byte[0] | . byte[1] | . byte[2] | . byte[3] | . byte[4] | . byte[5] | . byte[6] | . byte[7] | . byte[8] | . byte[9] | . byte |  | . byte[11] | . byte[12] | . byte[13] | . byte[14] | . byte[15] |
| nitble nitble | nibole. nible | nibte nible | nible. nibole | nibble) nible | nitble.nible | nioble nible | nibole. nitole | nibole. nible | nibule.nable | nibol e | nible | nibule.nible | nibut ef nibble | nitble nitble | nibule.nible | nibule nible |
| 04 | 812 | 1620 | 2428 | 3236 | 4044 | $48 \quad 52$ | 5660 | 6468 | 7276 | 80 | 84 | 8892 | 96100 | 104108 | 112116 | 120124 |

Figure 100.Vector Register elements

### 6.3.1 Vector Registers

There are 32 Vector Registers (VRs), each containing 128 bits. See Figure 101. All computations and other data manipulation are performed on data residing in Vector Registers, and results are placed into a VR.

| VR0 |  |
| :---: | :---: |
| VR1 |  |
| $\ldots$ |  |
| $\ldots$ |  |
| 0 | VR30 |

## Figure 101.Vector Registers

Depending on the instruction, the contents of a Vector Register are interpreted as a sequence of equal-length elements (bytes, halfwords, or words) or as a quadword. Each of the elements is aligned within the Vector Register, as shown in Figure 100. Many instructions perform a given operation in parallel on all elements in a Vector Register. Depending on the instruction, a byte, halfword, or word element can be interpreted as a signed-integer, an unsigned-integer, or a logical value; a word element can also be interpreted as a single-precision floating-point value. In the instruction descriptions, phrases like "signed-integer word element" are used as shorthand for "word element, interpreted as a signed-integer".

Load and Store instructions are provided that transfer a byte, halfword, word, or quadword between storage and a Vector Register.

### 6.3.2 Vector Status and Control Register

The Vector Status and Control Register (VSCR) is a special 32-bit register (not an SPR) that is read and written in a manner similar to the FPSCR in the Power ISA scalar floating-point unit. Special instructions (mfvscr and mtvscr) are provided to move the VSCR from and to a vector register. When moved to or from a vector register, the 32-bit VSCR is right justified in the 128 -bit vector register. When moved to a vector register, bits 0:95 of the vector register are cleared (set to 0 ).


Figure 102.Vector Status and Control Register
The bit definitions for the VSCR are as follows.

## Bit(s) Description

96:110 Reserved
$111 \quad$ Vector Non-Java Mode (NJ)
This bit controls how denormalized values are handled by Vector Floating-Point instructions.
0 Denormalized values are handled as specified by Java and the IEEE standard; see Section 6.6.1.
1 If an element in a source VR contains a denormalized value, the value 0 is used instead. If an instruction causes an Underflow Exception, the corresponding element in the target VR is set to 0 . In both cases the 0 has the same sign as the denormalized or underflowing value.

Reserved
Vector Saturation (SAT)

Every vector instruction having "Saturate" in its name implicitly sets this bit to 1 if any result of that instruction "saturates"; see Section 6.8. mtvscr can alter this bit explicitly. This bit is sticky; that is, once set to 1 it remains set to 1 until it is set to 0 by an mtvser instruction.

After the mfvscr instruction executes, the result in the target vector register will be architecturally precise. That is, it will reflect all updates to the SAT bit that could have been made by vector instructions logically preceding it in the program flow, and further, it will not reflect any SAT updates that may be made to it by vector instructions logically following it in the program flow. To implement this, processors may choose to make the mfvscr instruction execution serializing within the vector unit, meaning that it will stall vector instruction execution until all preceding vector instructions are complete and have updated the architectural machine state. This is permitted in order to simplify implementation of the sticky status bit (SAT) which would otherwise be difficult to implement in an out-of-order execution machine. The implication of this is that reading the VSCR can be much slower than typical Vector instructions, and therefore care must be taken in reading it, as advised in Section 6.5.1, to avoid performance problems.

The mtvscr is context synchronizing. This implies that all Vector instructions logically preceding an mtvscr in the program flow will execute in the architectural context ( NJ mode) that existed prior to completion of the mtvscr, and that all instructions logically following the mtvscr will execute in the new context ( NJ mode) established by the mtvscr.

### 6.3.3 VR Save Register

The VR Save Register (VRSAVE) is a 32-bit register in the fixed-point processor provided for application and operating system use; see Section 3.2.3.

## Programming Note

The VRSAVE register can be used to indicate which VRs are currently being used by a program. If this is done, the operating system could save only those VRs when an "interrupt" occurs (see Book III), and could restore only those VRs when resuming the interrupted program.

If this approach is taken it must be applied rigorously; if a program fails to indicate that a given VR is in use, software errors may occur that will be difficult to detect and correct because they are timing-dependent.

Some operating systems save and restore VRSAVE only for programs that also use other vector registers.

### 6.4 Vector Storage Access Operations

The Vector Storage Access instructions provide the means by which data can be copied from storage to a Vector Register or from a Vector Register to storage. Instructions are provided that access byte, halfword, word, and quadword storage operands. These instructions differ from the fixed-point and floating-point Storage Access instructions in that vector storage operands are assumed to be aligned, and vector storage accesses are performed as if the appropriate number of low-order bits of the specified effective address (EA) were zero. For example, the low-order bit of EA is ignored for halfword Vector Storage Access instructions, and the low-order four bits of EA are ignored for quadword Vector Storage Access instructions. The effect is to load or store the storage operand of the specified length that contains the byte addressed by EA.

If a storage operand is unaligned, additional instructions must be used to ensure that the operand is correctly placed in a Vector Register or in storage. Instructions are provided that shift and merge the contents of two Vector Registers, such that an unaligned quadword storage operand can be copied between storage and the Vector Registers in a relatively efficient manner.

As shown in Figure 100, the elements in Vector Registers are numbered; the high-order (or most significant) byte element is numbered 0 and the low-order (or least significant) byte element is numbered 15. The numbering affects the values that must be placed into the permute control vector for the Vector Permute instruction in order for that instruction to achieve the desired effects, as illustrated by the examples in the following subsections.

A vector quadword Load instruction for which the effective address (EA) is quadword-aligned places the byte in storage addressed by EA into byte element 0 of the target Vector Register, the byte in storage addressed by EA+1 into byte element 1 of the target Vector Register, etc. Similarly, a vector quadword Store instruction for which the EA is quadword-aligned places the contents of byte element 0 of the source Vector Register into the byte in storage addressed by EA, the contents of byte element 1 of the source Vector Register into the byte in storage addressed by $E A+1$, etc.

Figure 103 shows an aligned quadword in storage. Figure 104 shows the result of loading that quadword into a Vector Register or, equivalently, shows the contents that must be in a Vector Register if storing that Vector Register is to produce the storage contents shown in Figure 103.

When an aligned byte, halfword, or word storage operand is loaded into a Vector Register, the element (byte, halfword, or word respectively) that receives the data is the element that would have received the data had the entire aligned quadword containing the storage operand addressed by EA been loaded. Similarly, when a byte, halfword, or word element in a Vector Register is stored into an aligned storage operand (byte, halfword, or word respectively), the element selected to be stored is the element that would have been stored into the storage operand addressed by EA had the entire Vector Register been stored to the aligned quadword containing the storage operand addressed by EA. (Byte storage operands are always aligned.)

For aligned byte, halfword, and word storage operands, if the corresponding element number is known when the program is written, the appropriate Vector Splat and Vector Permute instructions can be used to copy or replicate the data contained in the storage operand after loading the operand into a Vector Register. An example of this is given in the Programming Note for Vector Splat; see page 259. Another example is to replicate the element across an entire Vector Register before storing it into an arbitrary aligned storage operand of the same length; the replication ensures that the correct data are stored regardless of the offset of the storage operand in its aligned quadword in storage.
00

10 | 00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09 | $0 A$ | $O B$ | $O C$ | $O D$ | $0 E$ | $0 F$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |  |

Figure 103.Aligned quadword storage operand

| 00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09 | $0 A$ | $0 B$ | $0 C$ | $0 D$ | $0 E$ | $0 F$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | $A$ | $B$ | $C$ | $D$ | E | F |

Figure 104.Vector Register contents for aligned quadword Load or Store

|  |  |  |  |  |  |  |  |  |  |  | 00 | 01 | 02 | 03 | 04 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | 05 | 06 | 07 | 08 | 09 | $0 A$ | $0 B$ | $0 C$ | $0 D$ | $0 E$ | $0 F$ |  |  |  |  |

Figure 105.Unaligned quadword storage operand

| Vhi |  |  |  |  |  |  |  |  |  |  |  | 00 | 01 | 02 | 03 | 04 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| VIo | 05 | 06 | 07 | 08 | 09 | OA | OB | OC | OD | OE | OF |  |  |  |  |  |
| Vt, Vs | 00 | 01 | 02 | 03 | 04 | 05 | 06 | 07 | 08 | 09 | OA | OB | OC | OD | OE | OF |

Figure 106.Vector Register contents

### 6.4.1 Accessing Unaligned Storage Operands

Figure 105 shows an unaligned quadword storage operand that spans two aligned quadwords. In the remainder of this section, the aligned quadword that contains the most significant bytes of the unaligned quadword is called the most significant quadword (MSQ) and the aligned quadword that contains the least significant bytes of the unaligned quadword is called the least significant quadword (LSQ). Because
the Vector Storage Access instructions ignore the low-order bits of the effective address, the unaligned quadword cannot be transferred between storage and a Vector Register using a single instruction. The remainder of this section gives examples of accessing unaligned quadword storage operands. Similar sequences can be used to access unaligned halfword and word storage operands.

## Programming Note

The sequence of instructions given below is one approach that can be used to load the unaligned quadword shown in Figure 105 into a Vector Register. In Figure 106 Vhi and Vlo are the Vector Registers that will receive the most significant quadword and least significant quadword respectively. VRT is the target Vector Register.

After the two quadwords have been loaded into Vhi and Vlo, using Load Vector Indexed instructions, the alignment is performed by shifting the 32-byte quantity Vhi || Vlo left by an amount determined by the address of the first byte of the desired data. The shifting is done using a Vector Permute instruction for which the permute control vector is generated by a Load Vector for Shift Left instruction. The Load Vector for Shift Left instruction uses the same address specification as the Load Vector Indexed instruction that loads the Vhi register; this is the address of the desired unaligned quadword.

The following sequence of instructions copies the unaligned quadword storage operand into register Vt.

```
\# Assumptions:
\# Rb \(!=0\) and contents of \(\mathrm{Rb}=0 \mathrm{xB}\)
lvx Vhi, 0, Rb \# load MSQ
lvsl \(\mathrm{Vp}, \mathrm{O}, \mathrm{Rb}\) \# set permute control vector
addi \(\mathrm{Rb}, \mathrm{Rb}, 16\) \# address of LSQ
lvx Vlo,0,Rb \# load LSQ
vperm Vt,Vhi,V1o,Vp \# align the data
```

The procedure for storing an unaligned quadword is essentially the reverse of the procedure for loading one. However, a read-modify-write sequence is required that inserts the source quadword into two aligned quadwords in storage. The quadword to be
stored is assumed to be in Vs; see Figure 106 The contents of Vs are shifted right and split into two parts, each of which is merged (using a Vector Select instruction) with the current contents of the two aligned quadwords (MSQ and LSQ) that will contain the most significant bytes and least significant bytes, respectively, of the unaligned quadword. The resulting two quadwords are stored using Store Vector Indexed instructions. A Load Vector for Shift Right instruction is used to generate the permute control vector that is used for the shifting. A single register is used for the "shifted" contents; this is possible because the "shifting" is done by means of a right rotation. The rotation is accomplished by specifying Vs for both components of the Vector Permute instruction. In addition, the same permute control vector is used on a sequence of 1 s and 0 s to generate the mask used by the Vector Select instructions that do the merging.

The following sequence of instructions copies the contents of Vs into an unaligned quadword in storage.

| \# Assumptions: |  |
| :---: | :---: |
| \# Rb ! $=0$ and contents of $\mathrm{Rb}=0 \mathrm{xB}$ |  |
| lvx | Vhi, $0, \mathrm{Rb}$ \# load current MSQ |
| lvsr | $\mathrm{Vp}, 0, \mathrm{Rb}$ \# set permute control vector |
| addi | $\mathrm{Rb}, \mathrm{Rb}, 16$ \# address of LSQ |
| lvx | Vlo, $0, \mathrm{Rb}$ \# load current LSQ |
| vspltisb | V1s,-1 \# generate the select mask bits |
| vspltisb | V0s, 0 |
| vperm | Vmask, V0s, V1s,Vp \# generate the select mask |
| vperm | Vs,Vs,Vs,Vp \# right rotate the data |
| vsel | Vlo,Vs,Vlo,Vmask \# insert LSQ component |
| vsel | Vhi,Vhi,Vs,Vmask \# insert MSQ component |
| stvx | Vlo,0,Rb \# store LSQ |
| addi | $\mathrm{Rb}, \mathrm{Rb},-16$ \# address of MSQ |
| stvx | Vhi, $0, \mathrm{Rb}$ \# store MSQ |

### 6.5 Vector Integer Operations

Many of the instructions that produce fixed-point integer results have the potential to compute a result value that cannot be represented in the target format. When this occurs, this unrepresentable intermediate value is converted to a representable result value using one of the following methods.

1. The high-order bits of the intermediate result that do not fit in the target format are discarded. This method is used by instructions having names that include the word "Modulo".
2. The intermediate result is converted to the nearest value that is representable in the target format (i.e., to the minimum or maximum representable value, as appropriate). This method is used by instructions having names that include the word "Saturate". An intermediate result that is forced to the minimum or maximum representable value as just described is said to "saturate".

An instruction for which an intermediate result saturates causes $\mathrm{VSCR}_{\text {SAT }}$ to be set to 1 ; see Section 6.3.2.
3. If the intermediate result includes non-zero fraction bits it is rounded up to the nearest fixed-point integer value. This method is used by the six Vector Average Integer instructions and by the Vector Multiply-High-Round-Add Signed Halfword Saturate instruction. The latter instruction then uses method 2 , if necessary.

## Programming Note

Because VSCR ${ }_{\text {SAT }}$ is sticky, it can be used to detect whether any instruction in a sequence of "Saturate"-type instructions produced an inexact result due to saturation. For example, the contents of the VSCR can be copied to a VR (mfvscr), bits other than the SAT bit can be cleared in the VR (vand with a constant), the result can be compared to zero setting CR6 (vcmpequb.), and a branch can be taken according to whether VSCR $_{\text {SAT }}$ was set to 1 (Branch Conditional that tests CR field 6).

Testing VSCR ${ }_{\text {SAT }}$ after each "Saturate"-type instruction would degrade performance considerably. Alternative techniques include the following:

- Retain sufficient information at "checkpoints" that the sequence of computations performed between one checkpoint and the next can be redone (more slowly) in a manner that detects exactly when saturation occurs. Test VSCR $_{\text {SAT }}$ only at checkpoints, or when redoing a sequence of computations that saturated.
- Perform intermediate computations using an element length sufficient to prevent saturation, and then use a Vector Pack Integer Saturate instruction to pack the final result to the desired length. (Vector Pack Integer Saturate causes results to saturate if necessary, and sets $\mathrm{VSCR}_{\text {SAT }}$ to 1 if any result saturates.)


### 6.5.1 Integer Saturation

Saturation occurs whenever the result of a saturating instruction does not fit in the result field. Unsigned saturation clamps results to zero (0) on underflow and to the maximum positive integer value ( $2^{\mathrm{n}}-1$, e.g. 255 for byte fields) on overflow. Signed saturation clamps results to the smallest representable negative number ( $-2^{n-1}$, e.g. -128 for byte fields) on underflow, and to the largest representable positive number ( $2^{n-1}-1$, e.g. +127 for byte fields) on overflow.

## Version 3.0

In most cases, the simple maximum/minimum saturation performed by the vector instructions is adequate. However, sometimes, e.g. in the creation of very high quality images, more complex saturation functions must be applied. To support this, the Vector facility provides a mechanism for detecting that saturation has occurred. The VSCR has a bit, the SAT bit, which is set to a one (1) anytime any field in a saturating instruction saturates. The SAT bit can only be cleared by explicitly writing zero to it. Thus SAT accumulates a summary result of any integer overflow or underflow that occurs on a saturating instruction.

Borderline cases that generate results equal to saturation values, for example unsigned $0+0=0$ and unsigned byte $1+254=255$, are not considered saturation conditions and do not cause SAT to be set.

The SAT bit can be set by the following types of instructions:

- Move To VSCR
- Vector Add Integer with Saturation
- Vector Subtract Integer with Saturation
- Vector Multiply-Add Integer with Saturation
- Vector Multiply-Sum with Saturation
- Vector Sum-Across with Saturation
- Vector Pack with Saturation
- Vector Convert to Fixed-point with Saturation

Note that only instructions that explicitly call for "saturation" can set SAT. "Modulo" integer instructions and floating-point arithmetic instructions never set SAT.

## Programming Note

The SAT state can be tested and used to alter program flow by moving the VSCR to a vector register (with mfvscr), then masking out bits 0:126 (to clear undefined and reserved bits) and performing a vector compare equal-to unsigned byte w/record (vcmpequb.) with zero to get a testable value into the condition register for consumption by a subsequent branch.

Since mfvscr will be slow compared to other Vector instructions, reading and testing SAT after each instruction would be prohibitively expensive. Therefore, software is advised to employ strategies that minimize checking SAT. For example: checking SAT periodically and backtracking to the last checkpoint to identify exactly which field in which instruction saturated; or, working in an element size sufficient to prevent any overflow or underflow during intermediate calculations, then packing down to the desired element size as the final operation (the vector pack instruction saturates the results and updates SAT when a loss of significance is detected).

### 6.6 Vector Floating-Point Operations

### 6.6.1 Floating-Point Overview

Unless $\mathrm{VSCR}_{\mathrm{NJ}}=1$ (see Section 6.3.2), the floating-point model provided by the Vector Facility conforms to The Java Language Specification (hereafter referred to as "Java"), which is a subset of the default environment specified by the IEEE standard (i.e., by ANSI/IEEE Standard 754-1985, "IEEE Standard for Binary Floating-Point Arithmetic"). For aspects of floating-point behavior that are not defined by Java but are defined by the IEEE standard, vector floating-point conforms to the IEEE standard. For aspects of floating-point behavior that are defined neither by Java nor by the IEEE standard but are defined by the "C9X Floating-Point Proposal" (hereafter referred to as "C9X"), vector floating-point conforms to C9X.

The single-precision floating-point data format, value representations, and computational models defined in Chapter 4. "Floating-Point Facility" on page 123 apply to vector floating-point except as follows.

- In general, no status bits are set to reflect the results of floating-point operations. The only exception is that $V_{S C R}$ SAT may be set by the Vector Convert To Fixed-Point Word instructions.
- With the exception of the two Vector Convert To Fixed-Point Word instructions and three of the four Vector Round to Floating-Point Integer instructions, all vector floating-point instructions that round use the rounding mode Round to Nearest.
- Floating-point exceptions (see Section 6.6.2) cannot cause the system error handler to be invoked.


## Programming Note

If a function is required that is specified by the IEEE standard, is not supported by the Vector Facility, and cannot be emulated satisfactorily using the functions that are supported by the Vector Facility, the functions provided by the Floating-Point Facility should be used; see Chapter 4.

### 6.6.2 Floating-Point Exceptions

The following floating-point exceptions may occur during execution of vector floating-point instructions.

- NaN Operand Exception
- Invalid Operation Exception
- Zero Divide Exception
- Log of Zero Exception
- Overflow Exception
- Underflow Exception

If an exception occurs, a result is placed into the corresponding target element as described in the following subsections. This result is the default result specified by Java, the IEEE standard, or C9X, as applicable.

Recall that denormalized source values are treated as if they were zero when $V S C R_{N J}=1$. This has the following consequences regarding exceptions.

- Exceptions that can be caused by a zero source value can be caused by a denormalized source value when $\mathrm{VSCR}_{\mathrm{NJ}}=1$.
- Exceptions that can be caused by a nonzero source value cannot be caused by a denormalized source value when $\mathrm{VSCR}_{\mathrm{NJ}}=1$.


### 6.6.2.1 NaN Operand Exception

A NaN Operand Exception occurs when a source value for any of the following instructions is a NaN.

- A vector instruction that would normally produce floating-point results
- Either of the two Vector Convert To Fixed-Point Word instructions
- Any of the four Vector Floating-Point Compare instructions

The following actions are taken:
If the vector instruction would normally produce floating-point results, the corresponding result is a source NaN selected as follows. In all cases, if the selected source NaN is a Signaling NaN it is converted to the corresponding Quiet NaN (by setting the high-order bit of the fraction field to 1) before being placed into the target element.
if the element in VRA is a NaN
then the result is that NaN
else if the element in VRB is a NaN
then the result is that NaN
else if the element in VRC is a NaN
then the result is that NaN
else if Invalid Operation exception
(Section 6.6.2.2)
then the result is the QNaN 0x7FCO_0000
If the instruction is either of the two Vector Convert To Fixed-Point Word instructions, the corresponding result is $0 \times 0000 \_0000 . V_{S C R}$ SAT is not affected.

If the instruction is Vector Compare Bounds Floating-Point, the corresponding result is 0xC000_0000.

If the instruction is one of the other Vector Floating-Point Compare instructions, the corresponding result is $0 \times 0000 \_0000$.

### 6.6.2.2 Invalid Operation Exception

An Invalid Operation Exception occurs when a source value or set of source values is invalid for the specified operation. The invalid operations are:

- Magnitude subtraction of infinities
- Multiplication of infinity by zero
- Reciprocal square root estimate of a negative, nonzero number or -infinity.
- Log base 2 estimate of a negative, nonzero number or -infinity.

The corresponding result is the QNaN 0x7FCO_0000.

### 6.6.2.3 Zero Divide Exception

A Zero Divide Exception occurs when a Vector Reciprocal Estimate Floating-Point or Vector Reciprocal Square Root Estimate Floating-Point instruction is executed with a source value of zero.

The corresponding result is an infinity, where the sign is the sign of the source value.

### 6.6.2.4 Log of Zero Exception

A Log of Zero Exception occurs when a Vector Log Base 2 Estimate Floating-Point instruction is executed with a source value of zero.

The corresponding result is -Infinity.

### 6.6.2.5 Overflow Exception

An Overflow Exception occurs under either of the following conditions.

- For a vector instruction that would normally produce floating-point results, the magnitude of what would have been the result if the exponent
range were unbounded exceeds that of the largest finite floating-point number for the target floating-point format.
- For either of the two Vector Convert To Fixed-Point Word instructions, either a source value is an infinity or the product of a source value and $2^{\text {UIM }}$ is a number too large in magnitude to be represented in the target fixed-point format.

The following actions are taken:

1. If the vector instruction would normally produce floating-point results, the corresponding result is an infinity, where the sign is the sign of the intermediate result.
2. If the instruction is Vector Convert To Unsigned Fixed-Point Word Saturate, the corresponding result is 0xFFFFF_FFFF if the source value is a positive number or +infinity, and is $0 \times 0000 \_0000$ if the source value is a negative number or -infinity. $\mathrm{VSCR}_{\mathrm{SAT}}$ is set to 1 .
3. If the instruction is Vector Convert To Signed Fixed-Point Word Saturate, the corresponding result is $0 x 7 F F F \_$FFFF if the source value is a positive number or +infinity., and is $0 \times 8000 \_0000$ if the source value is a negative number or -infinity. $\mathrm{VSCR}_{\mathrm{SAT}}$ is set to 1 .

### 6.6.2.6 Underflow Exception

An Underflow Exception can occur only for vector instructions that would normally produce floating-point results. It is detected before rounding. It occurs when a nonzero intermediate result computed as though both the precision and the exponent range were unbounded is less in magnitude than the smallest normalized floating-point number for the target floating-point format.

The following actions are taken:

1. If $\mathrm{VSCR}_{\mathrm{NJ}}=0$, the corresponding result is the value produced by denormalizing and rounding the intermediate result.
2. If $\mathrm{VSCR}_{\mathrm{NJ}}=1$, the corresponding result is a zero, where the sign is the sign of the intermediate result.

### 6.7 Vector Storage Access Instructions

The Vector Storage Access instructions compute the effective address (EA) of the storage to be accessed as described in Section 1.11.3, "Effective Address Calculation" on page 27. The low-order bits of the EA that would correspond to an unaligned storage operand are ignored.

The Load Vector Element Indexed and Store Vector Element Indexed instructions transfer a byte, halfword, or word element between storage and a Vector Register. The Load Vector Indexed and Store Vector Indexed instructions transfer an aligned quadword between storage and a Vector Register.

### 6.7.1 Storage Access Exceptions

Storage accesses will cause the system data storage error handler to be invoked if the program is not allowed to modify the target storage (Store only), or if the program attempts to access storage that is unavailable.

### 6.7.2 Vector Load Instructions

The aligned byte, halfword, word, or quadword in storage addressed by EA is loaded into register VRT.

## Programming Note

The Load Vector Element instructions load the specified element into the same location in the target register as the location into which it would be loaded using the Load Vector instruction.

## Load Vector Element Byte Indexed X-form

Ivebx VRT,RA,RB

| 31 | VRT | RA | RB |  | 7 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 61 |  |  |  |  |

$$
\begin{aligned}
& \text { if } R A=0 \text { then } b \leftarrow 0 \\
& \text { else } \quad \mathrm{b} \leftarrow(\mathrm{RA}) \\
& \mathrm{EA} \leftarrow \mathrm{~b}+(\mathrm{RB}) \\
& \text { eb } \leftarrow \mathrm{EA}_{60}: 63 \\
& \mathrm{VRT} \leftarrow \text { undefined } \\
& \text { if Big-Endian byte ordering then } \\
& \quad \mathrm{VRT}_{8 \times e \mathrm{eb}}: 8 \times \mathrm{eb}+7 \leftarrow \mathrm{MEM}(\mathrm{EA}, 1) \\
& \text { else } \\
& \quad \operatorname{VRT}_{120}-(8 \times \mathrm{eb}): 127-(8 \times \mathrm{eb}) \leftarrow \mathrm{MEM}(\mathrm{EA}, 1)
\end{aligned}
$$

Let the effective address (EA) be the sum $(R A \mid 0)+(R B)$.

Let eb be bits 60:63 of EA.
If Big-Endian byte ordering is used for the storage access, the contents of the byte in storage at address EA are placed into byte eb of register VRT. The remaining bytes in register VRT are set to undefined values.

If Little-Endian byte ordering is used for the storage access, the contents of the byte in storage at address EA are placed into byte $15-\mathrm{eb}$ of register VRT. The remaining bytes in register VRT are set to undefined values.

## Special Registers Altered:

None

## Load Vector Element Halfword Indexed X-form

| Ivehx VRT,RA,RB |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{rr}  & 31 \\ 0 & \\ \hline \end{array}$ | ${ }_{6}$ VRT | ${ }_{11} \mathrm{RA}$ | ${ }_{16} \mathrm{RB}$ | 21 | 39 | $\left\lvert\, \begin{aligned} & 1 \\ & 31\end{aligned}\right.$ |

$$
\begin{aligned}
& \text { if } R A=0 \text { then } b \leftarrow 0 \\
& \text { else } \quad b \leftarrow(R A) \\
& E A \leftarrow(b+(R B)) \& 0 x F F F F \_F F F F \_F F F F \_F F F E \\
& \text { eb } \leftarrow E A_{60: 63} \\
& \text { VRT } \leftarrow \text { undefined } \\
& \text { if Big-Endian byte ordering then } \\
& \quad V R T_{8 \times e b}: 8 \times e b+15 \leftarrow M E M(E A, 2) \\
& \text { else } \\
& \quad V_{R T} 112-(8 \times e b): 127-(8 \times e b) \leftarrow M E M(E A, 2)
\end{aligned}
$$

Let the effective address (EA) be the result of ANDing 0xFFFF_FFFF_FFFF_FFFE with the sum $(R A \mid 0)+(R B)$.

Let eb be bits 60:63 of EA.
If Big-Endian byte ordering is used for the storage access,

- the contents of the byte in storage at address EA are placed into byte eb of register VRT,
- the contents of the byte in storage at address $E A+1$ are placed into byte eb+1 of register VRT, and
- the remaining bytes in register VRT are set to undefined values.

If Little-Endian byte ordering is used for the storage access,

- the contents of the byte in storage at address EA are placed into byte $15-\mathrm{eb}$ of register VRT,
- the contents of the byte in storage at address $\mathrm{EA}+1$ are placed into byte 14 -eb of register VRT, and
- the remaining bytes in register VRT are set to undefined values.


## Special Registers Altered:

None

## Load Vector Element Word Indexed X-form

Ivewx VRT,RA,RB

| 31 | VRT |  | RA |  | RB |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  |  |  | 71 | 1 |  |

$$
\begin{aligned}
& \text { if } R A=0 \text { then } b \leftarrow 0 \\
& \text { else } \quad \mathrm{b} \leftarrow(\mathrm{RA}) \\
& \mathrm{EA} \leftarrow(\mathrm{~b}+(\mathrm{RB})) \& \text { OXFFFF_FFFF_FFFF_FFFC }^{\text {P }} \\
& \mathrm{eb} \leftarrow \mathrm{EA}_{60: 63} \\
& \text { VRT } \leftarrow \text { undefined } \\
& \text { if Big-Endian byte ordering then } \\
& \operatorname{VRT}_{8 \times e b}: 8 \times \mathrm{eb}+31 \leftarrow \operatorname{MEM}(E A, 4) \\
& \text { else } \\
& \operatorname{VRT}_{96-(8 \times e b)}: 127-(8 \times \mathrm{eb}) \leftarrow \operatorname{MEM}(E A, 4)
\end{aligned}
$$

Let the effective address (EA) be the result of ANDing OxFFFF_FFFF_FFFF_FFFC with the sum $(\mathrm{RA} \mid 0)+(\mathrm{RB})$.

Let eb be bits 60:63 of EA
If Big-Endian byte ordering is used for the storage access,

- the contents of the byte in storage at address EA are placed into byte eb of register VRT,
- the contents of the byte in storage at address $\mathrm{EA}+1$ are placed into byte eb+1 of register VRT,
- the contents of the byte in storage at address EA+2 are placed into byte eb+2 of register VRT,
- the contents of the byte in storage at address $E A+3$ are placed into byte eb+3 of register VRT, and
- the remaining bytes in register VRT are set to undefined values.

If if Little-Endian byte ordering is used for the storage access,

- the contents of the byte in storage at address EA are placed into byte $15-\mathrm{eb}$ of register VRT,
- the contents of the byte in storage at address $E A+1$ are placed into byte 14 -eb of register VRT,
- the contents of the byte in storage at address $\mathrm{EA}+2$ are placed into byte 13 -eb of register VRT,
- the contents of the byte in storage at address $\mathrm{EA}+3$ are placed into byte 12 -eb of register VRT, and
- the remaining bytes in register VRT are set to undefined values.


## Special Registers Altered:

None

## Load Vector Indexed X-form



| 31 | 6 | VRT | RA | RB |  | 103 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 11 |  |  |  |  |  |  |

$$
\begin{aligned}
& \text { if } R A=0 \text { then } b \leftarrow 0 \\
& \text { else } \quad b \leftarrow(\text { RA }) \\
& \text { EA } \leftarrow \mathrm{b}+(\mathrm{RB}) \\
& \text { VRT } \leftarrow \operatorname{MEM}\left(E A \& 0 x F F F F \_F F F F \_F F F F \_F F F 0, ~ 16\right)
\end{aligned}
$$

Let the effective address (EA) be the sum (RA|0)+(RB). The quadword in storage addressed by the result of EA ANDed with 0xFFFF_FFFF_FFFF_FFFO is loaded into VRT.

## Special Registers Altered:

None

## Load Vector Indexed Last X-form

IvxI VRT,RA,RB


$$
\begin{aligned}
& \text { if RA }=0 \text { then } b \leftarrow 0 \\
& \text { else } \quad b \leftarrow(R A) \\
& \text { EA } \leftarrow \mathrm{b}+(\mathrm{RB}) \\
& \text { VRT } \leftarrow \text { MEM }\left(E A \& 0 x F F F F \_F F F F \_F F F F \_F F F 0,16\right) \\
& \text { mark_as_not_likely_to_be_needed_again_anytime_soon (EA) }
\end{aligned}
$$

Let the effective address (EA) be the sum (RA|O)+(RB). The quadword in storage addressed by the result of EA ANDed with 0xFFFF_FFFF_FFFF_FFFO is loaded into VRT.

IvxI provides a hint that the quadword in storage addressed by EA will probably not be needed again by the program in the near future.

## Special Registers Altered:

 None
## Version 3.0

## Programming Note

On some implementations, the hint provided by the IvxI instruction and the corresponding hint provided by the stvxI, IvepxI, and stvepxI instructions are applied to the entire cache block containing the specified quadword. On such implementations, the effect of the hint may be to cause that cache block to be considered a likely candidate for replacement when space is needed in the cache for a new block. Thus, on such implementations, the hint should be used with caution if the cache block containing the quadword also contains data that may be needed by the program in the near future. Also, the hint may be used before the last reference in a sequence of references to the quadword if the subsequent references are likely to occur sufficiently soon that the cache block containing the quadword is not likely to be displaced from the cache before the last reference.

### 6.7.3 Vector Store Instructions

Some portion or all of the contents of VRS are stored into the aligned byte, halfword, word, or quadword in storage addressed by EA.

## Programming Note

The Store Vector Element instructions store the specified element into the same storage location as the location into which it would be stored using the Store Vector instruction.

## Store Vector Element Byte Indexed X-form

 stvebx VRS,RA,RB| 31 | VRS | RA | RB |  | 135 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 61 |  |  |  |  |

```
if \(R A=0\) then \(\mathrm{b} \leftarrow 0\)
else \(\quad \mathrm{b} \leftarrow(\mathrm{RA})\)
\(\mathrm{EA} \leftarrow \mathrm{b}+(\mathrm{RB})\)
\(\mathrm{eb} \leftarrow E A_{60: 63}\)
if Big-Endian byte ordering then
    \(\operatorname{MEM}(E A, 1) \leftarrow \operatorname{VRS}_{8 \times e b}: 8 \times \mathrm{eb}+7\)
else
    \(\left.\operatorname{MEM}(E A, 1) \leftarrow \operatorname{VRS}_{120-(8 \times e b}\right): 127-(8 \times e b)\)
```

Let the effective address (EA) be the sum (RA|O)+(RB).

Let eb be bits 60:63 of EA.
If Big-Endian byte ordering is used for the storage access, the contents of byte eb of register VRS are placed in the byte in storage at address EA.

If Little-Endian byte ordering is used for the storage access, the contents of byte $15-\mathrm{eb}$ of register VRS are placed in the byte in storage at address EA.

## Special Registers Altered:

None

## Programming Note

Unless bits 60:63 of the address are known to match the byte offset of the subject byte element in register VRS, software should use Vector Splat to splat the subject byte element before performing the store.

## Store Vector Element Halfword Indexed $X$-form

stvehx VRS,RA,RB

| 31 | VRS | RA | RB |  | 167 | 1 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |

$$
\begin{aligned}
& \text { if } R A=0 \text { then } b \leftarrow 0 \\
& \text { else } \quad b \leftarrow(R A) \\
& \text { EA } \leftarrow(b+(\mathrm{RB})) \& 0 \times F F F F \_F F F F_{-} F F F F \_F F F E \\
& \text { eb } \leftarrow \mathrm{EA}_{60}: 63 \\
& \text { if } \operatorname{Big}-E n d i a n ~ b y t e ~ o r d e r i n g ~ t h e n ~ \\
& \quad M E M(E A, 2) \leftarrow V R S_{8 \times e b}: 8 \times e b+15 \\
& \text { else } \\
& \quad M E M(E A, 2) \leftarrow \operatorname{VRS}_{112}-(8 \times \mathrm{eb}): 127-(8 \times \mathrm{eb})
\end{aligned}
$$

Let the effective address (EA) be the result of ANDing 0xFFFF_FFFF_FFFF_FFFE with the sum $(R A \mid 0)+(R B)$.

Let eb be bits 60:63 of EA.
If Big-Endian byte ordering is used for the storage access,

- the contents of byte eb of register VRS are placed in the byte in storage at address EA, and
- the contents of byte eb+1 of register VRS are placed in the byte in storage at address EA+1.

If Little-Endian byte ordering is used for the storage access,

- the contents of byte 15 -eb of register VRS are placed in the byte in storage at address EA, and
- the contents of byte 14-eb of register VRS are placed in the byte in storage at address EA+1.


## Special Registers Altered:

None

## Programming Note

Unless bits 60:62 of the address are known to match the halfword offset of the subject halfword element in register VRS software should use Vector Splat to splat the subject halfword element before performing the store.

## Store Vector Element Word Indexed X-form

stvewx VRS,RA,RB

| 31 | VRS | RA | RB |  | 199 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 61 |  |  |  |

```
if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(\textrm{RA}
EA \leftarrow(b + (RB))& OxFFFF_FFFF__FFFF_FFFC
eb}\leftarrowE\mp@subsup{A}{60:63}{
if Big-Endian byte ordering then
    MEM (EA,4) \leftarrowVRS }8\timese\textrm{eb}:8\timeseb+3
else
    MEM(EA,4)}\leftarrow\mp@subsup{\operatorname{VRS}}{96-(8\timeseb):127-(8\timeseb)}{
```

Let the effective address (EA) be the result of ANDing 0xFFFF_FFFF_FFFF_FFFC with the sum $(\mathrm{RA} \mid 0)+(\mathrm{RB})$.

Let eb be bits 60:63 of EA.
If Big-Endian byte ordering is used for the storage access,

- the contents of byte eb of register VRS are placed in the byte in storage at address EA,
- the contents of byte eb+1 of register VRS are placed in the byte in storage at address EA+1,
- the contents of byte eb+2 of register VRS are placed in the byte in storage at address EA+2, and
- the contents of byte eb+3 of register VRS are placed in the byte in storage at address EA+3.

If Little-Endian byte ordering is used for the storage access,

- the contents of byte $15-\mathrm{eb}$ of register VRS are placed in the byte in storage at address EA,
- the contents of byte 14 -eb of register VRS are placed in the byte in storage at address EA+1,
- the contents of byte 13-eb of register VRS are placed in the byte in storage at address EA+2, and
- the contents of byte 12-eb of register VRS are placed in the byte in storage at address EA+3.


## Special Registers Altered:

None

## Programming Note

Unless bits 60:61 of the address are known to match the word offset of the subject word element in register VRS, software should use Vector Splat to splat the subject word element before performing the store.

## Store Vector Indexed X-form

| Stvx |
| :--- |
| 31 VRS, RA, RB      <br> 0  6 RA RB  231 <br> 1       |

$$
\begin{aligned}
& \text { if } \mathrm{RA}=0 \text { then } \mathrm{b} \leftarrow 0 \\
& \text { else } \quad \mathrm{b} \leftarrow(\mathrm{RA}) \\
& \mathrm{EA} \leftarrow \mathrm{~b}+(\mathrm{RB}) \\
& \text { MEM }\left(E A \& 0 x F F F F \_F F F F \_F F F F \_F F F 0,16\right) \leftarrow(\mathrm{VRS})
\end{aligned}
$$

Let the effective address (EA) be the sum (RA|O)+(RB). The contents of VRS are stored into the quadword in storage addressed by the result of EA ANDed with 0xFFFF_FFFF_FFFF_FFFO.

## Special Registers Altered:

None

## Store Vector Indexed Last X-form

stvxI VRS,RA,RB


$$
\begin{aligned}
& \text { if } R A=0 \text { then } b \leftarrow 0 \\
& \text { else } \quad b \leftarrow(R A) \\
& \text { EA } \leftarrow \mathrm{b}+(\mathrm{RB}) \\
& \text { MEM }\left(E A \& 0 x F F F F \_F F F F \_F F F F \_F F F 0,16\right) \leftarrow \text { (VRS) } \\
& \text { mark_as_not_likely_to_be_needed_again_anytime_soon(EA) }
\end{aligned}
$$

Let the effective address (EA) be the sum $(R A \mid O)+(R B)$. The contents of VRS are stored into the quadword in storage addressed by the result of EA ANDed with 0xFFFF_FFFF_FFFF_FFFO.
stvxl provides a hint that the quadword in storage addressed by EA will probably not be needed again by the program in the near future.

## Special Registers Altered:

None

## Programming Note

See the Programming Note for the IvxI instruction on page 245 .

### 6.7.4 Vector Alignment Support Instructions

## Programming Note

The IvsI and Ivsr instructions can be used to create the permute control vector to be used by a subsequent vperm instruction (see page 262). Let $X$ and $Y$ be the contents of register VRA and VRB specified by the vperm. The control vector created by Ivsl causes the vperm to select the high-order 16 bytes of the result of shifting the 32-byte value $X$ II Y left by sh bytes. The control vector created by Ivsr causes the vperm to select the low-order 16 bytes of the result of shifting $X \| Y$ right by sh bytes.

## Programming Note

Examples of uses of IvsI, Ivsr, and vperm to load and store unaligned data are given in Section 6.4.1.

These instructions can also be used to rotate or shift the contents of a Vector Register left (Ivsl) or right (Ivsr) by sh bytes. For rotating, the Vector Register to be rotated should be specified as both register VRA and VRB for vperm. For shifting left, VRB for vperm should be a register containing all zeros and VRA should contain the value to be shifted, and vice versa for shifting right.

## Load Vector for Shift Left Indexed X-form

IvsI VRT,RA,RB

| 31 | 6 | VRT | RA | RB |  | 6 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 31 |  |  |  |  |  |  |

$$
\begin{aligned}
& \text { if } R A=0 \text { then } b \leftarrow 0 \\
& \text { else } \quad b \leftarrow(R A) \\
& \text { sh } \leftarrow(b+(R B))_{60: 63} \\
& \text { Switch(sh) }
\end{aligned}
$$

case ( $0 x 0$ ) : VRT $\leftarrow 0 x 000102030405060708090$ AOBOCODOEOF
case (0x1) : VRT $\leftarrow 0 x 0102030405060708090$ ABBOCODOEOF10
case (0x2) : VRT $\leftarrow 0 x 02030405060708090 A 0 B O C O D O E O F 1011$
case (0x3): VRT $\leftarrow 0 x 030405060708090$ AOBOCODOEOF101112
case (0x4) : VRT $\leftarrow 0 x 0405060708090$ AOBOCODOEOF10111213
case (0x5) : VRT $\leftarrow 0 x 05060708090$ AOBOCODOEOF1011121314
case (0x6) : VRT $\leftarrow 0 x 060708090 A 0 B O C O D O E O F 101112131415$
case (0x7) : VRT $\leftarrow 0 x 0708090 A 0 B 0 C O D O E O F 10111213141516$
case (0x8) : VRT $\leftarrow 0 x 08090$ AOBOCODOEOF1011121314151617
case (0x9) : VRT $\leftarrow 0 x 090 A 0 B O C O D O E O F 101112131415161718$ case (0xA) : VRT $\leftarrow 0 x 0 A O B O C O D O E O F 10111213141516171819$ case (0xB) : VRT $\leftarrow 0 x 0 B O C O D O E O F 101112131415161718191$ A case (0xC) : VRT $\leftarrow 0 x 0 C O D O E 0 F 101112131415161718191$ A1B case (0xD) : VRT $\leftarrow 0 x O D O E O F 101112131415161718191$ A1B1C case (0xE) : VRT $\leftarrow 0 x 0 E 0 F 101112131415161718191$ A1B1C1D case $(0 \mathrm{xF}):$ VRT $\leftarrow 0 \mathrm{x} 0 \mathrm{~F} 101112131415161718191 \mathrm{A1B1C1D1E}$

Let sh be bits $60: 63$ of the sum $(R A \mid 0)+(R B)$. Let $X$ be the 32 byte value $0 x 00$ || $0 x 01$ || $0 x 02$ || ... || 0x1E || $0 \times 1 F$.

Bytes sh to sh+15 of $X$ are placed into VRT.

## Special Registers Altered:

None

## Load Vector for Shift Right Indexed X-form

Ivsr
VRT,RA,RB

| 31 | VRT |  | RA | RB |  | 38 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  |  | 11 |  |  |  |

if $\mathrm{RA}=0$ then $\mathrm{b} \leftarrow 0$
else $\quad \mathrm{b} \leftarrow$ (RA)
$\mathrm{sh} \leftarrow(\mathrm{b}+(\mathrm{RB}))_{60: 63}$
switch (sh)
case (0x0) : VRT $\leftarrow 0 \times 101112131415161718191$ A1B1C1D1E1F
case (0x1) : VRT $\leftarrow 0 x 0 F 101112131415161718191$ A1B1C1D1E
case (0x2) : VRT $\leftarrow 0 \times 0$ OOF101112131415161718191A1B1C1D
case (0x3) : VRT $\leftarrow 0 x 0 D 0 E 0 F 101112131415161718191$ A1B1C
case (0x4) : VRT $\leftarrow 0 x 0 C O D O E 0 F 101112131415161718191$ A1B
case (0x5) : VRT $\leftarrow$ OxOBOCODOEOF101112131415161718191A
case (0x6) : VRT $\leftarrow 0 x 0 A O B O C O D O E O F 10111213141516171819$
case (0x7) : VRT $\leftarrow 0 \times 090$ AOBOCODOEOF101112131415161718
case (0x8) : VRT $\leftarrow 0 x 08090$ AOBOCODOEOF1011121314151617
case (0x9) : VRT $\leftarrow 0 x 0708090$ AOBOCODOEOF10111213141516
case (0xA) : VRT $\leftarrow 0 x 060708090$ AOBOCODOEOF101112131415
case (0xB) : VRT $\leftarrow 0 x 05060708090$ AOBOCODOEOF1011121314
case (0xC) : VRT $\leftarrow 0 x 0405060708090$ AOBOCODOEOF10111213
case (0xD) : VRT $\leftarrow 0 x 030405060708090$ AOBOCODOEOF101112
case (0xE) : VRT $\leftarrow 0 \times 02030405060708090$ AOBOCODOEOF1011
case (0xF) : VRT $\leftarrow 0 x 0102030405060708090 A 0 B O C O D O E O F 10$
Let sh be bits 60:63 of the sum (RA|0)+(RB). Let $X$ be the 32 -byte value $0 \times 00$ || $0 \times 01$ || $0 x 02$ || ... || 0x1E || $0 \times 1 \mathrm{~F}$.

Bytes 16 -sh to 31 -sh of $X$ are placed into VRT.

## Special Registers Altered:

None

### 6.8 Vector Permute and Formatting Instructions

### 6.8.1 Vector Pack and Unpack Instructions

## Vector Pack Pixel VX-form

| Vpkpx |
| :--- |
| 4 VRT, VRA, VRB     <br> 0  6 VRA VRB  <br> 11      |

```
do i = 0 to 63 by 16
    VR[VRT] i
    VR[VRT] 
    VR[VRT] }\mp@subsup{i}{i+6:i+10}{}\leftarrow\operatorname{VR}[VRA\mp@subsup{]}{i\times2+16:i\times2+20}{
    VR[VRT]}\mp@subsup{]}{i+11:i+15}{}\leftarrow\operatorname{VR}[VRA\mp@subsup{]}{i\times2+24:i\times2+28}{1
    VR[VRT] }\mp@subsup{\mp@code{i+64}}{}{~}\leftarrow\mathrm{ VR[VRB] }\mp@subsup{\mp@code{VN2+7}}{}{*
    VR[VRT] }\mp@subsup{\mp@code{i+65:i+69}}{}{~
    VR[VRT] }\mp@subsup{i}{i+70:i+74}{}\leftarrow\mathrm{ VR[VRB] i×2+16:i×2+20
    VR[VRT] }\mp@subsup{i}{i+75:i+79}{}\leftarrow\mathrm{ VR[VRB] i }\times2+24:i\times2+2
end
```

Let the source vector be the concatenation of the contents of VR[VRA] followed by the contents of VR[ VRB].

For each integer value i from 0 to 7 , do the following. Word element i in the source vector is packed to produce a 16 -bit value as described below.

- bit 7 of the first byte (bit 7 of the word)
- bits $0: 4$ of the second byte (bits $8: 12$ of the word)
- bits $0: 4$ of the third byte (bits $16: 20$ of the word)
- bits 0:4 of the fourth byte (bits 24:28 of the word)

The result is placed into halfword element i of VR[ VRT]

## Special Registers Altered:

None

## Programming Note

Each source word can be considered to be a 32-bit "pixel", consisting of four 8 -bit "channels". Each target halfword can be considered to be a 16-bit pixel, consisting of one 1-bit channel and three 5-bit channels. A channel can be used to specify the intensity of a particular color, such as red, green, or blue, or to provide other information needed by the application.

## Vector Pack Signed Doubleword Signed Saturate VX-form

```
vpksdss VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & & 1486 & \\
\hline 0 & & 6 & 11 & & & \\
\hline
\end{tabular}
```

```
src.qword[0] \leftarrow VR[VRA]
```

src.qword[0] \leftarrow VR[VRA]
src.qword[1] \leftarrowVR[VRB]
src.qword[1] \leftarrowVR[VRB]
do i = 0 to 3
do i = 0 to 3
VR[VRT].word[i] \leftarrowChop(Clamp( EXTS( src.dword[i]), -2 31,
VR[VRT].word[i] \leftarrowChop(Clamp( EXTS( src.dword[i]), -2 31,
2 31-1 ), 32 )
2 31-1 ), 32 )
end

```
end
```

Let doubleword elements 0 and 1 of $\operatorname{src}$ be the contents of VR[ VRA].

Let doubleword elements 2 and 3 of src be the contents of VR[ VRB].

For each integer value i from 0 to 3 , do the following.
The signed integer value in doubleword element $i$ of $s r c$ is placed into word element $i$ of VR[VRT] in signed integer format.

- If the value is greater than $2^{31 .} 1$ the result saturates to $2^{31} .1$.
- If the value is less than $\cdot 2^{31}$ the result saturates to- $2^{31}$.


## Special Registers Altered:

SAT

## Vector Pack Signed Doubleword Unsigned Saturate VX-form

vpksdus VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 1358 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  | 11 |  |

```
src.qword[0] \(\leftarrow\) VR[VRA]
src.qword[1] \(\leftarrow \operatorname{VR}[V R B]\)
do \(i=0\) to 3
    VR[VRT] .word[i] \(\leftarrow\) Chop (Clamp ( EXTS (src.dword[i]), 0, \(2^{32}-1\)
1, 32 )
end
```

Let doubleword elements 0 and 1 of $\operatorname{src}$ be the contents of VR[ VRA] .

Let doubleword elements 2 and 3 of $\operatorname{src}$ be the contents of VR[ VRB].

For each integer value i from 0 to 3 , do the following.
The signed integer value in doubleword element of $s r^{c}$ is placed into word element $i$ of VR[VRT] in unsigned integer format.

- If the value is greater than $2^{32}-1$ the result saturates to $2^{32}-1$.
- If the value is less than 0 the result saturates to 0 .


## Special Registers Altered:

SAT

## Vector Pack Signed Halfword Signed Saturate VX-form

vpkshss VRT,VRA,VRB

| 4 | VRT | VRA | VRB | 398 |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 611 |  |  |  |

```
do i=0 to 63 by 8
    Src1 \leftarrow EXTS((VRA) i\times2:i\times2+15)
    Src2 \leftarrow EXTS((VRB) i\times2:i\times2+15)
    VRT}\mp@subsup{\mp@code{i:i+7}}{}{\leftarrow}\leftarrow\mathrm{ Clamp(src1, -128, 127)}24:3
    VRT
end
```

Let the source vector be the concatenation of the contents of VRA followed by the contents of VRB.

For each integer value i from 0 to 15 , do the following.
Signed-integer halfword element $i$ in the source vector is converted to an signed-integer byte.

- If the value of the element is greater than 127 the result saturates to 127
- If the value of the element is less than -128 the result saturates to -128.

The low-order 8 bits of the result is placed into byte element $i$ of VRT.

Special Registers Altered:
SAT

## Vector Pack Signed Halfword Unsigned Saturate VX-form

vpkshus VRT,VRA,VRB

| 4 | VRT | VRA | VRB | 270 | 270 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |

```
do i=0 to 63 by }
    Src1}\leftarrow\operatorname{EXTS}((VRA) i\times2:i\times2+15
    SrC2 }\leftarrow\operatorname{EXTS}((VRB) i\times2:i\times2+15
    VRT
    VRT i+64:i+71
end
```

Let the source vector be the concatenation of the contents of VRA followed by the contents of VRB.

For each integer value i from 0 to 15 , do the following. Signed-integer halfword element i in the source vector is converted to an unsigned-integer byte.

- If the value of the element is greater than 255 the result saturates to 255
- If the value of the element is less than 0 the result saturates to 0 .

The low-order 8 bits of the result is placed into byte element $i$ of VRT.

Special Registers Altered: SAT

## Vector Pack Signed Word Signed Saturate VX-form

vpkswss VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  |  | 462 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  |  | 11 |  |  |  |

```
do i=0 to 63 by 16
    src1}\leftarrow\operatorname{EXTS}((VRA\mp@subsup{)}{i\times2:i\times2+31)}{
    Src2 \leftarrow EXTS ((VRB) i\times2:i\times2+31)
    VRT
    VRT
end
```

Let the source vector be the concatenation of the contents of VRA followed by the contents of VRB.

For each integer value i from 0 to 7 , do the following. Signed-integer word element $i$ in the source vector is converted to an signed-integer halfword.

- If the value of the element is greater than $2^{15}-1$ the result saturates to $2^{15}-1$
- If the value of the element is less than $-2^{15}$ the result saturates to $-2^{15}$.

The low-order 16 bits of the result is placed into halfword element $i$ of VRT.

Special Registers Altered: SAT

## Vector Pack Signed Word Unsigned Saturate VX-form

vpkswus VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 334 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 |  |  |

```
do i=0 to 63 by 16
```



```
    SrC2 }\leftarrow\operatorname{EXTS((VRB) i }\times2:i\times2+31
    VRT
    VRT}\mp@subsup{\textrm{i}+64:i+79}{}{~}\leftarrow\mathrm{ Clamp (src2, 0, 2'16-1) 16:31
```

Let the source vector be the concatenation of the contents of VRA followed by the contents of VRB.

For each integer value i from 0 to 7 , do the following.
Signed-integer word element $i$ in the source vector is converted to an unsigned-integer halfword.

- If the value of the element is greater than $2^{16}-1$ the result saturates to $2^{16}-1$
- If the value of the element is less than 0 the result saturates to 0 .

The low-order 16 bits of the result is placed into halfword element i of VRT.

```
Special Registers Altered:
    SAT
```


## Vector Pack Unsigned Doubleword Unsigned Modulo VX-form

vpkudum VRT,VRA,VRB

| 4 | 6 | VRT | VRA | VRB |  | 1102 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 16 |  |  |  |  |  |  |

```
if MSR.VEC then Vector_Unavailable()
src.qword[0] \leftarrow VR[VRA]
src.qword[1] \leftarrow VR[VRB]
do i = 0 to 3
    VR[VRT].word[i] \leftarrow Chop( EXTZ(src.dword[i]), 32 )
end
```

Let doubleword elements 0 and 1 of $\operatorname{src}$ be the contents of VR[ VRA].

Let doubleword elements 2 and 3 of $\operatorname{src}$ be the contents of VR[ VRB].

For each integer value i from 0 to 3 , do the following.
The contents of bits 32:63 of doubleword element $i$ of $s r c$ is placed into word element $i$ of VR[VRT].

## Vector Pack Unsigned Doubleword Unsigned Saturate VX-form

vpkudus VRT,VRA,VRB

| 4 | VRT |  | VRA | VRB | 1230 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  | 21 |

if MSR.VEC then Vector_Unavailable()

```
src.qword[0] \leftarrow VR[VRA]
src.qword[1] \leftarrowVR[VRB]
do i = 0 to 3
    VR[VRT].word[i] \leftarrow Chop(Clamp( EXTZ(src.dword[i]), 0, 2 22-1
), 32 )
end
```

Let doubleword elements 0 and 1 of $\operatorname{src}$ be the contents of VR[ VRA] .

Let doubleword elements 2 and 3 of $\operatorname{src}$ be the contents of VR[ VRB].

For each integer value i from 0 to 3 , do the following.
The unsigned integer value in doubleword element $i$ of $s r c$ is placed into word element $i$ of VR[ VRT] in unsigned integer format.

- If the value of the element is greater than $2^{32}-1$ the result saturates to $2^{32}-1$


## Special Registers Altered:

SAT

## Vector Pack Unsigned Halfiword Unsigned Modulo VX-form

vpkuhum VRT,VRA,VRB

| 4 | 6 | VRT | VRA | VRB |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 614 |  |  |  |

```
do i=0 to 63 by }
    VRTi:i+7 }\leftarrow(VRA\mp@subsup{)}{i\times2+8:i\times2+15}{
    VRT}\mp@subsup{T}{i+64:i+71}{}\leftarrow(VRB) i\times2+8:i\times2+1
```

end

Let the source vector be the concatenation of the contents of VRA followed by the contents of VRB.

For each integer value i from 0 to 15, do the following.
The contents of bits $8: 15$ of halfword element $i$ in the source vector is placed into byte element i of VRT.

## Special Registers Altered:

None

## Special Registers Altered:

None

## Vector Pack Unsigned Halfword Unsigned Saturate VX-form

vpkuhus VRT,VRA,VRB

| 4 | VRT | VRA | VRB | 142 |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |

```
do i=0 to 63 by }
    src1 \leftarrow EXTZ((VRA) i\times2:i\times2+15)
    Src2 \leftarrow EXTZ((VRB) 
    VRT
    VRT
end
```

Let the source vector be the concatenation of the contents of VRA followed by the contents of VRB.

For each integer value i from 0 to 15 , do the following. Unsigned-integer halfword element $i$ in the source vector is converted to an unsigned-integer byte.

- If the value of the element is greater than 255 the result saturates to 255 .

The low-order 8 bits of the result is placed into byte element i of VRT.

## Special Registers Altered:

 SAT
## Vector Pack Unsigned Word Unsigned Modulo VX-form

vpkuwum VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 78 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 |  |  |  |

```
do i=0 to 63 by 16
    VRT
    VRT
end
```

Let the source vector be the concatenation of the contents of VRA followed by the contents of VRB.

For each integer value ifrom 0 to 7 , do the following.
The contents of bits 16:31 of word element $i$ in the source vector is placed into halfword element $i$ of VRT.

## Special Registers Altered:

None

Vector Pack Unsigned Word Unsigned Saturate VX-form
vpkuwus VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 206 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 611 |  |  |  |  |

```
do i=0 to 63 by 16
    src1 \leftarrow EXTZ ((VRA) i\times2:i\times2+31 )
    Src2 \leftarrow EXTZ((VRB) i\times2:i\times2+31)
    VRT
    VRT i+64:i+79}\leftarrow\leftarrowClamp(\operatorname{src}2,0, 2 26-1 ) 16:31
end
```

Let the source vector be the concatenation of the contents of VRA followed by the contents of VRB.

For each integer value i from 0 to 7 , do the following. Unsigned-integer word element $i$ in the source vector is converted to an unsigned-integer halfword.

- If the value of the element is greater than $2^{16}-1$ the result saturates to $2^{16}-1$.

The low-order 16 bits of the result is placed into halfword element i of VRT.

## Special Registers Altered:

SAT

## Vector Unpack High Pixel VX-form

vupkhpx VRT,VRB

| 4 | VRT |  | III | VRB |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 646 |  |  |  |

$$
\begin{aligned}
& \text { do } i=0 \text { to } 63 \text { by } 16 \\
& \begin{array}{ll}
V R T_{i \times 2: i \times 2+7} & \leftarrow \operatorname{EXTS}\left((\mathrm{VRB})_{i}\right) \\
V R T_{i \times 2+8: i \times 2+15} & \leftarrow \operatorname{EXTZ}\left((\mathrm{VRB})_{i+1: i+5}\right) \\
V R T_{i \times 2+16: i \times 2+23} \leftarrow \operatorname{EXTZ}\left((\mathrm{VRB})_{i+6: i+10}\right) \\
V R T_{i \times 2+24: i \times 2+31} \leftarrow \operatorname{EXTZ}\left((\mathrm{VRB})_{i+11: i+15}\right)
\end{array}
\end{aligned}
$$

end
For each vector element $i$ from 0 to 3 , do the following. Halfword element $i$ in VRB is unpacked as follows.

- sign-extend bit 0 of the halfword to 8 bits
- zero-extend bits 1:5 of the halfword to 8 bits
- zero-extend bits 6:10 of the halfword to 8 bits
- zero-extend bits 11:15 of the halfword to 8 bits

The result is placed in word element $i$ of VRT.

## Special Registers Altered:

None

## Programming Note

The source and target elements can be considered to be 16-bit and 32 -bit "pixels" respectively, having the formats described in the Programming Note for the Vector Pack Pixel instruction on page 250.

## Programming Note

Notice that the unpacking done by the Vector Unpack Pixel instructions does not reverse the packing done by the Vector Pack Pixel instruction. Specifically, if a 16 -bit pixel is unpacked to a 32 -bit pixel which is then packed to a 16 -bit pixel, the resulting 16-bit pixel will not, in general, be equal to the original 16-bit pixel (because, for each channel except the first, Vector Unpack Pixel inserts high-order bits while Vector Pack Pixel discards low-order bits).

## Vector Unpack Low Pixel VX-form

vupklpx VRT,VRB

| 04 | $6$ | ${ }^{\prime \prime} \quad \text { III }$ | ${ }_{16} \text { VRB }$ | 21 | 974 | 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |

$$
\begin{aligned}
& \text { do } \mathrm{i}=0 \text { to } 63 \text { by } 16
\end{aligned}
$$

$$
\begin{aligned}
& \operatorname{VRT}_{1 \times 2+16: 1 \times 2+23} \leftarrow \operatorname{EXTZ}\left((\mathrm{VRB})_{i+70: i+74}\right) \\
& \operatorname{VRT}_{1 \times 2 \times 24: 1 \times 2+31} \leftarrow \operatorname{EXTZ}\left((\operatorname{VRB}){ }_{i+775: i+79}\right)
\end{aligned}
$$

end
For each vector element $i$ from 0 to 3, do the following. Halfword element $i+4$ in VRB is unpacked as follows.

- sign-extend bit 0 of the halfword to 8 bits
- zero-extend bits 1:5 of the halfword to 8 bits
- zero-extend bits 6:10 of the halfword to 8 bits
- zero-extend bits 11:15 of the halfword to 8 bits

The result is placed in word element $i$ of VRT.

## Special Registers Altered: <br> None

Vector Unpack High Signed Byte VX-form
Vopkhsb

| 4 | VRT | VRB |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | III | VRB |  |
| 16 | 526 |  |  |  |  |

$$
\begin{aligned}
& \text { do } i=0 \text { to } 63 \text { by } 8 \\
& \operatorname{VRT}_{i \times 2: i \times 2+15} \leftarrow \operatorname{EXTS}\left((\mathrm{VRB})_{i: i+7}\right) \\
& \text { end }
\end{aligned}
$$

For each vector element ifrom 0 to 7 , do the following. Signed-integer byte element in VRB is sign-extended to produce a signed-integer halfword and placed into halfword element $i$ in VRT.

Special Registers Altered:
None

## Vector Unpack High Signed Halfword VX-form

Vupkhsh

| 4 | VRT,VRB |  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | III | VRB |  |
| 11 |  | 16 | 21 |  | 390 |  |

$$
\begin{aligned}
& \text { do } i=0 \text { to } 63 \text { by } 16 \\
& \text { VRT }_{i \times 2: i \times 2+31} \leftarrow \operatorname{EXTS}\left((\mathrm{VRB})_{i: i+15}\right) \\
& \text { end }
\end{aligned}
$$

For each vector element $i$ from 0 to 3 , do the following. Signed-integer halfword element $i$ in VRB is sign-extended to produce a signed-integer word and placed into word element $i$ in VRT.

## Special Registers Altered: <br> None

## Vector Unpack High Signed Word VX-form

vupkhsw VRT,VRB

| 4 | VRT |  | III | VRB | 1614 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 61 |  |  |  |  |

$\operatorname{VR}[V R T] . d w o r d[0] \leftarrow C h o p(\operatorname{EXTS}(V R[V R B] \cdot w o r d[0]), 64)$
VR[VRT].$d w o r d[1] \leftarrow C h o p(\operatorname{EXTS}(\operatorname{VR}[V R B] \cdot w o r d[1]), 64)$

For each integer value i from 0 to 1 , do the following. The signed integer value in word element $i$ of VR[VRB] is sign-extended and placed into doubleword element $i$ of VR[ VRT].

## Special Registers Altered:

None

Vector Unpack Low Signed Byte VX-form vupklsb VRT,VRB

| 4 | VRT |  | III | VRB |  | 654 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 |  |  |  |

```
do i=0 to 63 by }
    VRT i\times2:i\times2+15}<<\operatorname{EXTS}((VRB) i+64:i+71
end
```

For each vector element $i$ from 0 to 7 , do the following. Signed-integer byte element i+8 in VRB is sign-extended to produce a signed-integer halfword and placed into halfword element $i$ in VRT.

Special Registers Altered:
None

## Vector Unpack Low Signed Halfword VX-form

vupklsh VRT,VRB

| 4 | VRT |  | III | VRB |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 718 |  |  |

$$
\begin{aligned}
& \text { do } i=0 \text { to } 63 \text { by } 16 \\
& \text { VRT }_{i \times 2: i \times 2+31} \leftarrow \operatorname{EXTS}\left((\mathrm{VRB})_{i+64: i+79}\right) \\
& \text { end }
\end{aligned}
$$

For each vector element $i$ from 0 to 3 , do the following. Signed-integer halfword element $i+4$ in VRB is sign-extended to produce a signed-integer word and placed into word element $i$ in VRT.

## Special Registers Altered:

None

## Vector Unpack Low Signed Word VX-form

vupklsw VRT,VRB

| 4 | VRT |  | III |  | VRB |  | 1742 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |  |

```
VR[VRT]. dword[0] \(\leftarrow\) Chop( EXTS (VR[VRB]. word[2]), 64 )
\(\operatorname{VR}[\operatorname{VRT}] . d w o r d[1] \leftarrow C h o p(\operatorname{EXTS}(V R[V R B] . \operatorname{word}[3]), 64)\)
```

For each integer value i from 0 to 1 , do the following. The signed integer value in word element $\mathrm{i}+2$ of VR[VRB] is sign-extended and placed into doubleword element $i$ of VR[ VRT].

Special Registers Altered:
None

### 6.8.2 Vector Merge Instructions

## Vector Merge High Byte VX-form

Vmrghb

| 4 | VRT, VRA, VRB |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | VRA | VRB | 12 |
| 16 |  | 21 |  | 31 |

```
do i=0 to 63 by }
    VRT i 
    VRT i\times2+8:i\times2+15}<<(VRB) i:i+7
end
```

For each vector element $i$ from 0 to 7 , do the following. Byte element $i$ in VRA is placed into byte element $2 \times i$ in VRT.

Byte element $i$ in VRB is placed into byte element $2 \times i+1$ in VRT.

Special Registers Altered:
None

Vector Merge High Halfword VX-form
vmrghh VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 76 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  |  |  |

$$
\begin{aligned}
& \text { do } i=0 \text { to } 63 \text { by } 16 \\
& \qquad V R T_{i \times 2: i \times 2+15}^{\leftarrow}(\mathrm{VRA})_{i: i+15} \\
& \qquad \operatorname{VRT}_{i \times 2+16: i \times 2+31} \leftarrow(\mathrm{VRB})_{i: i+15} \\
& \text { end }
\end{aligned}
$$

For each vector element $i$ from 0 to 3 , do the following. Halfword element $i$ in VRA is placed into halfword element $2 \times i$ in VRT.

Halfword element $i$ in VRB is placed into halfword element $2 \times i+1$ in VRT.

Special Registers Altered:
None

## Vector Merge Low Byte VX-form

vmrglb VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 268 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 611 |  |  |  |  |

```
do i=0 to 63 by }
    VRT i\times2:i\times2+7}<\leftarrow(VRA) i+64:i+71
    VRT
end
```

For each vector element $i$ from 0 to 7 , do the following. Byte element i+8 in VRA is placed into byte element $2 \times i$ in VRT.

Byte element $\mathrm{i}+8$ in VRB is placed into byte element $2 \times i+1$ in VRT.

## Special Registers Altered:

None

Vector Merge Low Halfword VX-form
vmrglh VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 332 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  | 11 |  |  |

```
do i=0 to 63 by 16
    VRT i\times2:i\times2+15}\leftarrow\leftarrow(VRA) i+64:i+79
    VRT
end
```

For each vector element $i$ from 0 to 3 , do the following. Halfword element $i+4$ in VRA is placed into halfword element $2 \times \mathrm{i}$ in VRT.

Halfword element $i+4$ in VRB is placed into halfword element $2 \times i+1$ in VRT.

## Special Registers Altered:

None

Vector Merge High Word VX-form
vmrghw VRT,VRA,VRB

| 4 | VRT | VRA | VRB | 140 |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 61 |  |  |  |

> do $\quad i=0$ to 63 by 32
> $\operatorname{VRT}_{i \times 2: i \times 2+31} \leftarrow(\mathrm{VRA})_{i: i+31}$
> $\operatorname{VRT}_{i \times 2+32: i \times 2+63} \leftarrow(\mathrm{VRB})_{i: i+31}$
end
For each vector element $i$ from 0 to 1 , do the following. Word element $i$ in VRA is placed into word element $2 \times i$ in VRT.

Word element $i$ in VRB is placed into word element $2 \times i+1$ in VRT.

The word elements in the high-order half of VRA are placed, in the same order, into the even-numbered word elements of VRT. The word elements in the high-order half of VRB are placed, in the same order, into the odd-numbered word elements of VRT.

## Special Registers Altered:

None

## Vector Merge Low Word VX-form

$$
\text { vmrglw } \quad \text { VRT,VRA,VRB }
$$

| 4 | VRT | VRA | VRB | 396 |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  |  |  |  |

> do $i=0$ to 63 by 32 $$
V R T_{i \times 2: i \times 2+31} \leftarrow(\mathrm{VRA})_{i+64: i+95}
$$ $\quad \mathrm{VRT}_{i \times 2+32: i \times 2+63} \leftarrow(\mathrm{VRB})_{i+64: i+95}$ end

For each vector element $i$ from 0 to 1 , do the following. Word element i+2 in VRA is placed into word element $2 \times i$ in VRT.

Word element i+2 in VRB is placed into word element $2 \times i+1$ in VRT.

Special Registers Altered:
None

## Vector Merge Even Word VX-form

vmrgew VRT,VRA,VRB

| 4 | VRT | VRA | VRB | 1932 |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 61 |  |  |  |

> if MSR. VEC=O then Vector Unavailable(l)
> VR[VRT]. word[0] $\leftarrow$ VR[VRā]. word[0]
> VR[VRT]. wor d[1] $\leftarrow$ VR[VRB], word[0]
> VR[VRT]. word[2] $\leftarrow$ VR[VRA], word[2]
> VR[VRT]. word[3] $\leftarrow$ VR[VRB]. word[2]

The contents of word element 0 of VR[VRA] are placed into word element 0 of VR[ VRT].

The contents of word element 0 of VR[VRB] are placed into word element 1 of VR[ VRT].

The contents of word element 2 of VR[VRA] are placed into word element 2 of VR[ VRT].

The contents of word element 2 of VR[VRB] are placed into word element 3 of VR[ VRT].
vmrgew is treated as a Vector instruction in terms of resource availability.

## Special Registers Altered

None

## Vector Merge Odd Word VX-form

vmrgow VRT,VRA,VRB

| 4 | VRT | VRA | VRB | 1676 |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 61 |  |  |  |  |

$$
\begin{aligned}
& \text { if MSR. VEC=O then Vector_Unavailablel) } \\
& \text { VR[VRT], word[0] } \leftarrow \text { VR[VRĀ], word[1] } \\
& \text { VR[ VRT]. wor d[1] } \leftarrow \text { VR[ VRB]. wor d[1] } \\
& \text { VR[ VRT]. word[2] } \leftarrow \text { VR[ VRA]. word [ } 3 \text { ] } \\
& \text { VR[ VRT]. word[3] } \leftarrow \text { VR[ VRB]. wor d [3] }
\end{aligned}
$$

The contents of word element 1 of VR[ VRA] are placed into word element 0 of V R[ VRT].

The contents of word element 1 of VR[ VRB] are placed into word element 1 of VR[VRT].

The contents of word element 3 of VR[VRA] are placed into word element 2 of VR[VRT].

The contents of word element 3 of VR[VRB] are placed into word element 3 of V R[ VRT] .
vmrgow is treated as a Vector instruction in terms of resource availability.

Special Registers Altered
None

### 6.8.3 Vector Splat Instructions

## Programming Note

The Vector Splat instructions can be used in preparation for performing arithmetic for which one source vector is to consist of elements that all have the same value (e.g., multiplying all elements of a Vector Register by a constant).

## Vector Splat Byte VX-form

| vspltb VRT,VRB, UIM |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 4 | $6$ |  UIM <br> 11 12 | 16 VRB | 21 | 524 | 31 |

$$
\begin{aligned}
& \mathrm{b} \leftarrow \text { UIM } \| 0 \mathrm{~b} 000 \\
& \text { do } \mathrm{i}=0 \text { to } 127 \text { by } 8 \\
& \operatorname{VRT}_{\mathrm{i}: \mathrm{i}+7} \leftarrow(\mathrm{VRB})_{\mathrm{b}: \mathrm{b}+7} \\
& \text { end }
\end{aligned}
$$

For each integer value i from 0 to 15 , do the following.
The contents of byte element UIM in VRB are placed into byte element $i$ of VRT.

## Special Registers Altered:

None

## Vector Splat Halfword VX-form

vsplth VRT,VRB,UIM

| 4 | VRT | I/ | UIM | VRB |  | 588 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 611 |  |  |  |  |

$$
\begin{aligned}
& \mathrm{b} \leftarrow \text { UIM } \| 0 \mathrm{~b} 0000 \\
& \text { do } \mathrm{i}=0 \text { to } 127 \text { by } 16 \\
& \operatorname{VRT}_{\mathrm{i}: \mathrm{i}+15} \leftarrow(\mathrm{VRB})_{\mathrm{b}: \mathrm{b}+15} \\
& \text { end }
\end{aligned}
$$

For each integer value i from 0 to 7 , do the following. The contents of halfword element UIM in VRB are placed into halfword element i of VRT.

## Special Registers Altered:

None

## Vector Splat Word VX-form

Vspltw

| 4 | VRT, VRB, UIM |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | VRT |

$\mathrm{b} \leftarrow$ UIM || 0 b00000
do $\mathrm{i}=0$ to 127 by 32

$$
\operatorname{VRT}_{i: i+31} \leftarrow\left(\mathrm{VRB}_{b: b+31}\right.
$$

end
For each integer value i from 0 to 3 , do the following.
The contents of word element UIM in VRB are placed into word element $i$ of VRT.

## Special Registers Altered:

None

## Vector Splat Immediate Signed Byte <br> VX-form

```
vspltisb VRT,SIM
```

| 4 | VRT |  | SIM | I/I |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 680 | 78 |  |  |  |

```
do i=0 to 127 by 8
    VRT
end
```

For each integer value i from 0 to 15 , do the following.
The value of the SIM field, sign-extended to 8 bits, is placed into byte element i of VRT.

Special Registers Altered:
None

Vector Splat Immediate Signed Halfword VX-form
vspltish VRT,SIM

| 4 | VRT |  | SIM | III |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  |  | 844 |  |  |  |

```
do i=0 to 127 by 16
    VRT
end
```

For each integer value i from 0 to 7 , do the following.
The value of the SIM field, sign-extended to 16 bits, is placed into halfword element i of VRT.

Special Registers Altered:
None

## Vector Splat Immediate Signed Word

VX-form
vspltisw VRT,SIM

| 4 | VRT |  | SIM | III |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 11 |  | 908 |  |  |  |  |

```
do i=0 to 127 by 32
    VRT i:i+31}\leftarrow\leftarrow\operatorname{EXTS}(SIM, 32
end
```

For each vector element $i$ from 0 to 3, do the following.
The value of the SIM field, sign-extended to 32 bits, is placed into word element $i$ of VRT.

None

### 6.8.4 Vector Permute Instruction

The Vector Permute instruction allows any byte in two source Vector Registers to be copied to any byte in the target Vector Register. The bytes in a third source Vector Register specify from which byte in the first two source Vector Registers the corresponding target byte is to be copied. The contents of the third source Vector Register are sometimes referred to as the "permute control vector".

## Vector Permute VA-form

VRerm

| 4 | VRT, VRA, VRB, VRC |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | VRA | VRB | VRC |  |
| 11 | 43 |  |  |  |  |  |

```
if MSR.VEC=O then Vector_Unavailablell
scc.qword[0] & Vr[VRa]
scc.qword[1] &VR[VRB]
do i = 0 to 15
    index & VR[VRC], byte[i], bit[3:7]
    VR[VRT], byte[i] & src.byte[index]
end
```

Let the source vector be the concatenation of the contents of VR[VRA] followed by the contents of VR[ VRB].

For each integer value i from 0 to 15 , do the following. Let index be the value specified by bits 3:7 of byte element $i$ of VR[VRC].

The contents of byte element index of sic are placed into byte element $i$ of VR[VRT].

## Special Registers Altered:

 None
## Vector Permute Right-indexed VA-form

vpermr VRT,VRA,VRB,VRC

| 4 | VRT | VRA | VRB | VRC |  | 59 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 61 |  |  |  |  |

if MSR. VEC=O then Vector_Unavailablel)

```
scc.qword[0] & VR[VRA]
scc.qword[1]}\leftarrowVR[VRB
do i = 0 to 15
    index & VR[VRC]. byte[i], bit[3:7]
    Vr[VRT], byte[i] & scc, byte[31.| ndex]
end
```

Let the source vector be the concatenation of the contents of VR[VRA] followed by the contents of VR[ VRB].

For each integer value i from 0 to 15 , do the following. Let index be the value specified by bits 3:7 of byte element $i$ of VR[ VRC].

The contents of byte element 31 -index of src are placed into byte element i of VR[ VRT].

Special Registers Altered:
None

### 6.8.5 Vector Select Instruction

## Vector Select VA-form

```
vsel VRT,VRA,VRB,VRC
```

| 4 | VRT | VRA | VRB | VRC | 42 | 42 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  |  |  |  |

```
if MSR.VEC=O then Vector_Unavailable()
do i = 0 to 127
    mask \leftarrowVR[VRC], bit[i]
    VR[VRT],bit[i] \leftarrow(mask=0) ? VR[VRA],bit[i]: VR[VRB], bit[i]
end
```

For each bit in VR[VRC] that contains the value 0, the corresponding bit in VR[VRA] is placed into the corresponding bit of VR[VRT]. Otherwise, the corresponding bit in VR[VRB] is placed into the corresponding bit of VR[ VRT] .

## Special Registers Altered:

None

## Version 3.0

### 6.8.6 Vector Shift Instructions

The Vector Shift instructions rotate or shift the contents of a Vector Register or a pair of Vector Registers left or right by a specified number of bytes (vslo, vsro, vsldoi) or bits (vsl, vsr). Depending on the instruction, this "shift count" is specified either by the contents of a Vector Register or by an immediate field in the instruction. In the former case, 7 bits of the shift count register give the shift count in bits ( $0 \leq$ count $\leq 127$ ). Of these 7 bits, the high-order 4 bits give the number of complete bytes by which to shift and are used by vslo and vsro; the low-order 3 bits give the number of remaining bits by which to shift and are used by vsl and vsr.

## Programming Note

A pair of these instructions, specifying the same shift count register, can be used to shift the contents of a Vector Register left or right by the number of bits (0-127) specified in the shift count register. The following example shifts the contents of register Vx left by the number of bits specified in register $V y$ and places the result into register $V z$.

```
vslo Vz,Vx,Vy
vsl Vz,Vz,Vy
```


## Vector Shift Left Double by Octet Immediate VA-form

| VRIdoi |
| :--- |
| 4 VRT, VRA, VRB, SHB       <br> 0  6 VRA VRB 1 SHB  <br> 12        |

if MSR.VEC=O then Vector_Unavailablel)
src.qword[0] $\leftarrow$ VR[VRA]
scc.aword [1] $\leftarrow$ VR[ VRB]
VR[VRT] \& src. byte[SHB: SHB+15]
Let the source vector be the concatenation of the contents of VR[VRA] followed by the contents of VR[VRB]. Bytes SHB:SHB+15 of the source vector are placed into VR[ VRT].

## Special Registers Altered:

None

## Vector Shift Left VX-form

VsI

| 4 | VRT, VRA, VRB |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | VRA | VRB |  |
| 11 |  |  |  |  |  |

```
if MSR.VEC=O then Vector_Unavailable()
shb &VR[VRB], bit[125:127]<<3
t\leftarrow1
do i =0 to 15
    t tt & (VR[VNB], byte[i], bit[5:7]=sh)
end
if t=1 then
    VR[VRT] &VR[VRA] << sh
else
    VR[VRT] \leftarrowundefi ned
```

The contents of VR[VRA] are shifted left by the number of bits specified in bits 125:127 of VR[ VRB].

- Bits shifted out of bit 0 are lost.
- Zeros are supplied to the vacated bits on the right.

The result is place into VR[VRT], except if, for any byte element in register VR[ VRB] , the low-order 3 bits are not equal to the shift amount, then VR[VRT] is undefined.

## Special Registers Altered:

None

## Vector Shift Left by Octet VX-form

Vslo

| 4 | VRT, VRA, VRB |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | VRA | VRB |
| 16 | 1036 |  |  |  |  |

## if MSR.VEC=0 then Vector_Unavailable()

shb $\leftarrow$ VR[VRB], bit $[121: 124] \ll 3$

$$
\text { VR[ VRT] } \leftarrow V R[V R A] \ll \operatorname{shb}
$$

The contents of VR[VRA] are shifted left by the number of bytes specified in bits 121:124 of VR[ VRB].

- Bytes shifted out of byte 0 are lost.
- Zeros are supplied to the vacated bytes on the right.

The result is placed into VR[ VRT].

## Special Registers Altered:

None

## Vector Shift Right VX-form

| Vsr |
| :--- |
| 4 VRT, VRA, VRB     <br> 0  6  VRA VRB <br> 11      |

```
if MSR.VEC=O then Vector_Unavailable|
sh}\leftarrow\mathrm{ VR[ VVB], bit[ [25: 127]
t\leftarrow1
do i = 0 to 15
    t \leftarrowt & (VR[VRB], byte[i],bit[5:7]=sh)
end
if t=1 then
    VR[VRT] & VR[VRA] >> sh
else
    VR[VRT] & undefined
```

The contents of VR[ VRA] are shifted right by the number of bits specified in bits 125:127 of VR[ VRB].

- Bits shifted out of bit 127 are lost.
- Zeros are supplied to the vacated bits on the left.

The result is place into VR[ VRT], except if, for any byte element in register VR[ VRB] , the low-order 3 bits are not equal to the shift amount, then VR[ VRT] is undefined.

## Special Registers Altered:

None

## Vector Shift Right by Octet VX-form

Vsro

| 4 | VRT, VRA, VRB |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | VRT | VRA | VRB |
| 11 |  | 1100 |  |  |  |

```
if MSR. VEC=O then Vector_Unavailablel)
shb \(\leftarrow \operatorname{VR[VRB],~bit[121:124]\ll 3~}\)
\[
\text { VR[VRT] } \operatorname{VR[VRA]~\gg shb~}
\]
```

The contents of VR[VRA] are shifted right by the number of bytes specified in bits 121:124 of V R[ VRB].

- Bytes shifted out of byte 15 are lost.
- Zeros are supplied to the vacated bytes on the left.

The result is placed into VR[ VRT].

## Special Registers Altered:

None

## Programming Note

A double-register shift by a dynamically specified number of bits (0-127) can be performed in six instructions. The following example shifts Vw \|| Vx left by the number of bits specified in Vy and places the high-order 128 bits of the result into Vz .

| vslo | Vt 1, Vw, Vy | \# shift high-order reg left |
| :---: | :---: | :---: |
| vsl | Vt $1, \mathrm{Vt}$ 1, Vy |  |
| vsububm | Vt 3, Vo, Vy | \# adjust shift count ( (V) $=0$ ) |
| vsio | Vt $2, V \mathrm{~V}, \mathrm{Vt} 3$ | \# shift low-order reg right |
| vsi | Vt2, Vt 2, Vt 3 |  |
| vor | Vz,Vt 1, Vt 2 | \# merge to get final result |

## Vector Shift Left Variable VX-form

vslv VRT,VRA,VRB

| 4 | VRT | VRA | VRB | 1860 |  | 31 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |

```
if MSR.VEC=O then Vector_Unavailable_Interrupt()
src.byte[0:15]}\leftarrowVR[VRA
src.byte[16]}\leftarrow0\times0
do i = 0 to 15
    sh}\leftarrowVR[VRB], byte[i],bit[5:7
    VR[VRT],byte[i] & src.byte[i:i +1], bit[sh:sh+7]
end
```

Let bytes $0: 15$ of $s i c$ be the contents of VR[ VRA]. Let byte 16 of src be the value $0 \times 00$.

For each integer value i from 0 to 15, do the following. Let $s h$ be the value in bits 5:7 of byte element $i$ of VR[ VRB].

The contents of bits $s h: s h+7$ of the halfword in byte elements $\mathrm{i}: \mathrm{i}+1$ of sic are placed into byte element $i$ of VR[ VRT].

Special Registers Altered:
None

Vector Shift Right Variable VX-form
Vsrv

| 4 | VRT,VRA, VRB |  |  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | VRA | VRB |  | 1796 |  |
| 16 |  |  |  |  |  |  |  |

if MSR.VEC=O then Vector_Unavailable_Interrupt()
src.byte[0] $\leftarrow 0 x 00$
src. byte[1:16] $\leftarrow$ VR[VRA]
do $i=0$ to 15
$s h \leftarrow V R[V R B]$, byte[i], bit[5:7]
VR[VRT], byte[i] \&src.byte[i:ill].bit[8•sh:15•sh]
end
Let bytes $0: 15$ of $s r c$ be the contents of VR[ VRA].
Let bytes 16 of src be the value $0 \times 00$.
For each integer value i from 0 to 15, do the following. Let sh be the value in bits 5:7 of byte element i of VR[ VRB]

The contents of bits $8 \cdot 5 \mathrm{~h}: 15 \cdot \mathrm{sh}$ of the halfword in byte elements $i: i+1$ of $\operatorname{sic}$ are placed into byte element i of VR[ VRT].

## Special Registers Altered:

None

## Programming Note

Assume vskl contains a vector of packed 7－bit values，A located in bits $0: 6, B$ located in bits $7: 13, C$ located in bits 14：20，etc．．

| \＃VSRC | оbaaaaaaab，оbвbвbвbcC，obcccccodo，obdoddeeee， |
| :---: | :---: |
| \＃ | obeEefffff，ObFFgGGGGG，ObGHHHHHH，ObIIIIIIII， |
| \＃ |  |
| $\#$ | ObN⿱㇒⿴囗⿱一一口儿 |

Assume the following registers are pre－loaded as follows，


The leftmost seven packed 7－bit values can be unpacked into byte elements 0 to 6 using vsrv with vSHCNT1．


The next seven packed 7－bit values can then be unpacked into byte elements 7 to 13 using vsrv with vSHCNT2．

| Vsiv | VTMP2，VTMP1，VSHCT2 | \＃VTMP2 | $=\{$ ObOAAAAAAA， | ObABBBBBBB， | ObBCCCCCCC， |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | \＃ | ObDEEEEEEE， | ObEFFFFFFF， | ObFGGGGGGG， |  |
|  |  | \＃ | ObHIIIIIII， | ，OblJJJJJノ」， | Ob］KKKKKKK， |  |
|  |  | \＃ | ObLMMMMMMM， | ObMNNNNNNN， | Ob0000000p， |  |

The next two packed 7－bit values can then be unpacked into byte elements 14 to 15 using vsrv with vSHCNT3．

vsiv vTMP3，vTMP2，vSHCT3 | $\#$ vTMP $=\left\{\begin{array}{l}\text { ObOAAAAAAA，ObABBBBBBB，ObBCCCCCCC，ObCDDDDDDD，} \\ \\ \\ \#\end{array}\right.$ | ObDEEEEEE，ObEFFFFFFF，ObFGGGGGGG，ObGHHHHHH， |  |
| ---: | :--- | :--- |
|  | $\#$ | ObHIIIIIII，ObIJJJJJJ，Ob］KKKKKKK，ObKLLLLLLL， |
|  | $\#$ | ObLMMMMMM，ObMNNNNNN，ObNOOOOOOO，ObOPPPPPPP $\} ;$ |

The most－significant bit in each byte element is masked off to produce a vector of sixteen unsigned byte elements．


The vector of sixteen unsigned byte elements can be further unpacked to two vectors of eight unsigned halfword elements using a vupkhsb and a vupklsb．

```
vupkhsb vTMP5, vTMP4 # vTMP5 ={0bOOOOOOOO_OAAAAAAA,ObOOOOOOOO_OBBBBBBB,\ldots };
vupklsb vTMP6, vTMP4 # vTMP6 ={0b00000000_01IIIIII,0b00000000_0|||J||, ...};
```

The resultant two vectors of eight unsigned halfword elements can then be further unpacked to four vectors of four unsigned word elements using two vupkhsh and two vupklsh instructions．

| vupkhsh | vRESULTO，vTMP5 | \＃VRESULTO $=\left\{0600000000 \_00000000 \_00000000 \_\right.$OAAAAAAA， |
| :---: | :---: | :---: |
| vupklsh | vRESULT1，VTMP5 | \＃vRESULT1 $=\{0600000000$＿00000000＿00000000＿OEEEEEE， |
| vupkhsh | vRESULT2，vTMP6 | \＃vRESULT2 $=\{0600000000$＿00000000＿00000000＿01111111， |
| vupklsh | vRESULT3，vTMP6 | \＃VRESULT3 $=\{0600000000$＿00000000＿00000000＿OMMMMMMM， |

### 6.8.7 Vector Extract Element Instructions

## Vector Extract Unsigned Byte VX-form

vextractub VRT,VRB,UIM

| 4 | VRT | 1 <br> 11 | UIM | VRB |  | 525 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |

```
if MSR.vec=0 then Vector_Unavaliable()
src & VR[ VRB] , byte[ Ul M]
VR[vRT].dword[0] & ExTz64(src)
VR[ VRT].dwor d[1] & 0x0000_0000_0000_0000
```

The contents of byte element $U I M$ of VR[ VRB] are placed into bits 56:63 of VR[VRT]. The contents of the remaining byte elements of VR[ VRT] are set to 0 .

## Special Registers Altered: <br> None

## Vector Extract Unsigned Halfword VX-form

| vextractuh |
| :--- |
| 4 VRT, VRB, UIM      <br> 0  6 11 UIM VRB  <br> 16       |

if MSR. VEC=0 then Vector Unavailable()
SrC $\leftarrow$ VR[VRB] , byte e[ UI M: Ul M+1]
VR[VRT]. dword[0] $\leftarrow$ Extz64(Sic)
VR[VRT]. dwor d[1] © Ox0000_0000_0000_0000
The contents of byte elements UIM:UIM+1 of VR[VRB] are placed into halfword element 3 of VR[VRT]. The contents of the remaining halfword elements of VR[ VRT] are set to 0 .

If the value of $U I M$ is greater than 14 , the results are undefined.

Special Registers Altered:
None

Vector Extract Doubleword VX-form
vextractd VRT,VRB,UIM

| 4 | VRT | $/$ | UIM | VRB |  | 717 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 | 12 | 16 | 21 |

if MSR.VEC=0 then Vector_Unavailable()
sic \& VR[VRB]. byte[ [U M: UI M+7]
VR[ VRT]. dwor dil O] $\leftarrow$ Sic
VR[VRT]. dwor d[1] $\leftarrow$ Ox0000_0000_0000_0000
The contents of byte elements UI M: UI M+7 of VR[VRB] are placed into VR[VRT]. The contents of doubleword element 1 of VR[ VRT] are set to 0 .

If the value of $U I M$ is greater than 8 , the results are undefined.

Special Registers Altered:
None

### 6.8.8 Vector Insert Element Instructions

Vector Insert Byte VX-form
VRT, VRB,UIM
vinsertb

| 4 | VRT |  | 11 | UIM | VRB |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 781 | 31 |  |  |

if MSR. VEC=O then Vector_Unavailable()
VR[VRT] , byte[ UI M] $\leftarrow$ VR[ VRB], byte[7]
The contents of byte element 1 of VR[VRB] are placed into byte element UIM of VR[VRT]. The contents of the remaining byte elements of VR[ VRT] are not modified.

## Special Registers Altered:

None

## Vector Insert Halfword VX-form

vinserth VRT,VRB,UIM

| 4 | 6 | VRT |  | UIM | VRB |  | 845 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 12 |  |  |  |

if MSR.VEC=O then Vector_Unavailablell
VR[VRT], byte[Ul M: Ul M+1] $\leftarrow$ VR[VRB], hwor d[3]
The contents of halfword element 3 of VR[VRB] are placed into byte elements $U \| M: U I M+1$ of VR[VRT]. The contents of the remaining byte elements of VR[VRT] are not modified.

If the value of $U I M$ is greater than 14 , the results are undefined.

## Special Registers Altered:

None

Vector Insert Doubleword VX-form
vinsertd VRT,VRB,UIM

| 04 | ${ }_{6}$ VRT |  | ${ }_{16}$ VRB | 21 | 973 | 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |

if MSR.VEC=O then Vector_Unavailablel)
VR[VRT], byte[Ul M: Ul M+7] $\leftarrow$ VR[VRB]. dwor d[0]
The contents of doubleword element 0 of VR[VRB] are placed into byte elements UIM: UIM+7 of VR[VRT]. The contents of the remaining byte elements of VR[VRT] are not modified.

If the value of $U I M$ is greater than 8 , the results are undefined.

## Special Registers Altered:

None

### 6.9 Vector Integer Instructions

### 6.9.1 Vector Integer Arithmetic Instructions

### 6.9.1.1 Vector Integer Add Instructions

## Vector Add and Write Carry-Out Unsigned Word VX-form

vaddcuw VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 384 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |

```
do i=0 to 127 by 32
    aop}\leftarrow\operatorname{EXTZ((VRA) i:i+31)
    bop }\leftarrow\operatorname{EXTZ((VRB) i:i+31)
    VRT i:i+31}\mp@code{\leftarrowChop( ( aop + +int bop ) >> ui 32,1)
end
```

For each integer value i from 0 to 3 , do the following. Unsigned-integer word element $i$ in VRA is added to unsigned-integer word element i in VRB. The carry out of the 32 -bit sum is zero-extended to 32 bits and placed into word element $i$ of VRT.

## Special Registers Altered:

None

Vector Add Signed Byte Saturate VX-form
vaddsbs
VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 768 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |

```
do i=0 to 127 by 8
    aop}\leftarrow\operatorname{EXTS}(VR\mp@subsup{A}{i:i+7}{}
    bop }\leftarrow\operatorname{EXTS}(VR\mp@subsup{B}{i:i+7}{\prime}
    VRT i:i+7}<\leftarrow\mathrm{ Clamp( aop + int bop, -128, 127 ) 24:31
end
```

For each integer value i from 0 to 15 , do the following. Signed-integer byte element $i$ in VRA is added to signed-integer byte element i in VRB.

- If the sum is greater than 127 the result saturates to 127.
- If the sum is less than -128 the result saturates to -128.

The low-order 8 bits of the result are placed into byte element i of VRT.

Special Registers Altered:
SAT

## Vector Add Signed Halfword Saturate VX-form

vaddshs VRT,VRA,VRB

| 4 | VRT |  | VRA | VRB |  | 832 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  |  |  |  |

```
do i=0 to 127 by 16
    aop}\leftarrow\operatorname{EXTS}((VRA\mp@subsup{)}{i:i+15}{
    bop}\leftarrow\operatorname{EXTS}((VRB) i:i+15
    VRT
end
```

For each integer value i from 0 to 7 , do the following. Signed-integer halfword element $i$ in VRA is added to signed-integer halfword element i in VRB.

- If the sum is greater than $2^{15}-1$ the result saturates to $2^{15}-1$
- If the sum is less than $-2^{15}$ the result saturates to $-2^{15}$.

The low-order 16 bits of the result are placed into halfword element i of VRT.

## Special Registers Altered:

SAT

## Vector Add Signed Word Saturate VX-form

vaddsws VRT,VRA,VRB

| 4 | VRT | VRA | VRB | 896 |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |

```
do i=0 to 127 by 32
    aop}\leftarrow\operatorname{EXTS}((VRA) i:i+31
    bop}\leftarrow\operatorname{EXTS}((VRB) i:i+31
    VRT}\mp@subsup{\mp@code{i:i+31}}{}{~}\operatorname{Clamp(aop +int bop, -2 21, 231-1)
end
```

For each integer value i from 0 to 3 , do the following. Signed-integer word element $i$ in VRA is added to signed-integer word element i in VRB.

- If the sum is greater than $2^{31}-1$ the result saturates to $2^{31}-1$.
- If the sum is less than $-2^{31}$ the result saturates to $-2^{31}$.

The low-order 32 bits of the result are placed into word element i of VRT.

## Special Registers Altered: SAT

## Vector Add Unsigned Byte Modulo VX-form

```
vaddubm VRT,VRA,VRB
```

| 4 | VRT | VRA | VRB |  | 0 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 |  |  |

For each integer value i from 0 to 15 , do the following.
Unsigned-integer byte element $i$ in VRA is added to unsigned-integer byte element i in VRB.

The low-order 8 bits of the result are placed into byte element i of VRT.

## Special Registers Altered:

## None

## Programming Note

vaddubm can be used for unsigned or signed-integers.

```
do i=0 to 127 by 8
```

do i=0 to 127 by 8
aop}\leftarrow\operatorname{EXTZ}((VRA) i:i+7
aop}\leftarrow\operatorname{EXTZ}((VRA) i:i+7
bop }\leftarrow\operatorname{EXTZ ((VRB) i:i+7 )
bop }\leftarrow\operatorname{EXTZ ((VRB) i:i+7 )
VRT i:i+7}\mp@code{\leftarrowChop( aop + +int bop, 8)
VRT i:i+7}\mp@code{\leftarrowChop( aop + +int bop, 8)
end

```
end
```

syicu-ticge.

$$
-
$$

## Vector Add Unsigned Doubleword Modulo VX-form

vaddudm VRT,VRA,VRB

| 4 | VRT |  | VRA | VRB |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  |  | 192 |  |  |  |

```
do i = 0 to 1
    aop \leftarrow VR[VRA].dword[i]
    bop }\leftarrow\textrm{VR[VRB].dword[i]
    VR[VRT].dword[i] \leftarrow Chop( aop + int bop, 64 )
end
```

For each integer value i from 0 to 1 , do the following. The integer value in doubleword element i of VR[VRB] is added to the integer value in doubleword element $i$ of VR[VRA].

The low-order 64 bits of the result are placed into doubleword element $i$ of VR[ VRT].

## Special Registers Altered:

None

## Programming Note

vaddudm can be used for signed or unsigned integers.

## Vector Add Unsigned Halfword Modulo VX-form

vadduhm VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 64 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  | 11 |  |

```
do i=0 to 127 by 16
    aop }\leftarrow\operatorname{EXTZ}((VRA) i:i+15
    bop \leftarrow EXTZ((VRB) i:i+15)
    VRT}\mp@subsup{T}{:i+15}{}\leftarrow\mathrm{ Chop( aop +int bop, 16)
end
```

For each integer value i from 0 to 7 , do the following. Unsigned-integer halfword element $i$ in VRA is added to unsigned-integer halfword element i in VRB.

The low-order 16 bits of the result are placed into halfword element i of VRT.

Special Registers Altered:
None

Programming Note
vadduhm can be used for unsigned or signed-integers.

## Vector Add Unsigned Word Modulo VX-form

vadduwm VRT,VRA,VRB

| 4 | VRT |  | VRA | VRB | 128 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |

do $i=0$ to 127 by 32
aop $\leftarrow \operatorname{EXTZ}\left((\mathrm{VRA})_{i: i+31}\right)$
bop $\leftarrow \operatorname{EXTZ}\left((\mathrm{VRB})_{i: i+31}\right)$
temp $\leftarrow$ aop $+_{\text {int }}$ bop
$\operatorname{VRT}_{\mathrm{i}: i+31} \leftarrow$ Chop $\left(\right.$ aop $+_{\text {int }}$ bop, 32$)$
end
For each integer value i from 0 to 3 , do the following. Unsigned-integer word element $i$ in VRA is added to unsigned-integer word element i in VRB.

The low-order 32 bits of the result are placed into word element i of VRT.

Special Registers Altered: None

Programming Note
vadduwm can be used for unsigned or signed-integers.

## Vector Add Unsigned Byte Saturate VX-form

## vaddubs VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 512 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |

```
do i=0 to 127 by 8
    aop }\leftarrow\operatorname{EXTZ}((VRA\mp@subsup{)}{i:i+7}{}
    bop}\leftarrow\operatorname{EXTZ}((VRB\mp@subsup{)}{i:i+7}{}
    VRT i:i+7
end
```

For each integer value i from 0 to 15 , do the following. Unsigned-integer byte element $i$ in VRA is added to unsigned-integer byte element i in VRB.

- If the sum is greater than 255 the result saturates to 255 .

The low-order 8 bits of the result are placed into byte element i of VRT.

Special Registers Altered:
SAT

## Vector Add Unsigned Halfword Saturate VX-form

```
vadduhs VRT,VRA,VRB
```

| 4 | VRT | VRA | VRB |  | 576 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 | 11 |  |  |

```
do i=0 to 127 by 16
    aop}\leftarrow\operatorname{EXTZ}((VRA) i:i+15
    bop}\leftarrow\operatorname{EXTZ}((VRB) i:i+15
    VRT
end
```

For each integer value i from 0 to 7 , do the following. Unsigned-integer halfword element $i$ in VRA is added to unsigned-integer halfword element i in VRB.

- If the sum is greater than $2^{16}-1$ the result saturates to $2^{16}-1$.

The low-order 16 bits of the result are placed into halfword element $i$ of VRT.

## Special Registers Altered:

SAT

## Vector Add Unsigned Word Saturate VX-form

vadduws VRT,VRA,VRB

| 4 | VRT | VRA | VRB |  | 640 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 61 |  |  |  |  |

```
do i=0 to 127 by 32
    aop \leftarrow EXTZ((VRA) i:i+31)
    bop \leftarrow EXTZ((VRB) i:i+31)
    VRT}\mp@subsup{T}{i:i+31}{}\leftarrowClamp(aop +int bop, 0, 232-1
end
```

For each integer value i from 0 to 3 , do the following. Unsigned-integer word element i in VRA is added to unsigned-integer word element i in VRB.

- If the sum is greater than $2^{32}-1$ the result saturates to $2^{32}-1$.

The low-order 32 bits of the result are placed into word element i of VRT.

Special Registers Altered:
SAT

## Vector Add Unsigned Quadword Modulo VX-form

vadduqm VRT,VRA,VRB

| 4 | VRT | VRA | VRB | 256 |  |
| :---: | :---: | :---: | :---: | :---: | :---: |

if MSR.VEC=0 then Vector_Unavailablel)

```
srcl \leftarrowVR[VRA]
sic2 \leftarrowVR[VRB]
sum &EXTZ(srcl) + EXTZ(src2)
VR[VRT] \leftarrowChop(sum, 128)
```

Let srcl be the integer value in VR[ VRA].
Let $\mathrm{src}_{2}$ be the integer value in VR[ VRB].
srcl and srcl can be signed or unsigned integers.
The rightmost 128 bits of the sum of srcl and srcl are placed into VR[ VRT].

## Special Registers Altered:

None

## Vector Add Extended Unsigned Quadword Modulo VA-form

vaddeuqm VRT,VRA,VRB,VRC

| 4 | VRT | VRA | VRB | VRC |  | 60 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  | 21 |  |

```
if MSR.VEC=0 then Vector_Unavailablel)
sicl &VR[VRA]
src2 \leftarrowVR[VRB]
cin}\leftarrowVR[VRC],bit[127
sum \leftarrowEXTZ(srcl) + EXTZ(src2) + EXTZ(cin)
VR[VRT] \leftarrowChop(sum, 128)
```

Let srcl be the integer value in VR[ VRA].
Let $\operatorname{src} 2$ be the integer value in VR[ VRB].
Let cin be the integer value in bit 127 of VR[ VRC].
srcl and srcl can be signed or unsigned integers.
The rightmost 128 bits of the sum of $\mathrm{srcl}, \mathrm{src}$, and cin are placed into VR[ VRT].

## Special Registers Altered:

None

## Vector Add \& write Carry Unsigned Quadword VX-form

```
vaddcuq VRT,VRA,VRB
```

| 4 | VRT | VRA | VRB |  | 320 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 11 |  | 6 |  |  |  |

```
if MSR.VEC=O then Vector_Unavailable()
```

```
srcl \leftarrowVR[VRA]
```

srcl \leftarrowVR[VRA]
src2 \leftarrowVR[VRB]
src2 \leftarrowVR[VRB]
sum \& EXTZ(src1) + EXTZ(src2)
sum \& EXTZ(src1) + EXTZ(src2)
VR[VRT] \leftarrowChop( EXTZ( Chop(sum>>128, 1) ), 128 )

```

Let srcl be the integer value in VR[ VRA].
Let src2 be the integer value in VR[ VRB].
srcl and srcl can be signed or unsigned integers.
The carry out of the sum of srcl and srcl is placed into VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Add Extended \& write Carry Unsigned Quadword VA-form}
vaddecuq VRT,VRA,VRB,VRC
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 4 & VRT & VRA & VRB & VRC & & 61 & \\
\hline 0 & & 6 & & 11 & & & 21 \\
\hline
\end{tabular}
```

if MSR.VEC=O then Vector_Unavailable()
srcl \leftarrowVR[VRA]
srC2}\leftarrowVR[VRB
cin \&VR[VRC],bit[127]
sum \leftarrowEXTZ(srcl) + EXTZ(src2) + EXTZ(cin)
VR[VRT] \leftarrowChop( EXTZ( Chop(sum >> 128, 1) ), 128)

```

Let srcl be the integer value in VR[ VRA].
Let \(\operatorname{src} 2\) be the integer value in VR[ VRB].
Let cin be the integer value in bit 127 of VR[ VRC].
srcl and srcl can be signed or unsigned integers.
The carry out of the sum of srcl, src2, and cin are placed into VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Version 3.0}

\section*{Programming Note}

The Vector Add Unsigned Quadword instructions support efficient wide-integer addition. The following code sequence can be used to implement a 512-bit signed or unsigned add operation.
\begin{tabular}{|c|c|c|}
\hline vadduqm & vS3, vA3, vB3 & \# bits 384:511 of sum \\
\hline vaddcua & vC3, vA3, vB3 & \# carry out of bit 384 of sum \\
\hline vaddeuqm & vS2, vA2, vB2, vC3 & \# bits 256:383 of sum \\
\hline vaddecuq & vC2, vA2, vB2, vC3 & \# carry out of bit 256 of sum \\
\hline vaddeuqm & vS1, vA1, vB1, vC2 & \# bits 128:255 of sum \\
\hline vaddecua & vC1, vA1, vB1, vC2 & \# carry out of bit 128 of sum \\
\hline vaddeugm & vSO, vAO, vBO, vC1 & \# bits 0:127 of sum \\
\hline
\end{tabular}

\subsection*{6.9.1.2 Vector Integer Subtract Instructions}

\section*{Vector Subtract and Write Carry-Out \\ Unsigned Word VX-form}
```

vsubcuw
VRT,VRA,VRB

```
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 1408 \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
aop }\leftarrow\operatorname{EXTZ}((VRA\mp@subsup{)}{i:i+31}{}
bop}\leftarrow\operatorname{EXTZ((VRB) i:i+31)
temp }\leftarrow(\mathrm{ aop + int ᄀbop +int 1) >> 32
VRT
end

```

For each integer value i from 0 to 3 , do the following. Unsigned-integer word element i in VRB is subtracted from unsigned-integer word element i in VRA. The complement of the borrow out of bit 0 of the 32-bit difference is zero-extended to 32 bits and placed into word element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Subtract Signed Byte Saturate VX-form}
vsubsbs
VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & VRT & VRA & VRB & & 1792 \\
\hline 11 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 8
aop}\leftarrow\operatorname{EXTS}((VRA\mp@subsup{)}{i:i+7}{}
bop}\leftarrow\operatorname{EXTS}((VRB\mp@subsup{)}{i:i+7}{}
VRT i:i+7}<<<Clamp(aop + int \negbop + int 1, -128, 127) 24:31
end

```

For each integer value i from 0 to 15 , do the following. Signed-integer byte element i in VRB is subtracted from signed-integer byte element \(i\) in VRA.
- If the intermediate result is greater than 127 the result saturates to 127 .
- If the intermediate result is less than -128 the result saturates to -128.

The low-order 8 bits of the result are placed into byte element i of VRT.

Special Registers Altered:
SAT

\section*{Vector Subtract Signed Halfword Saturate VX-form}
```

vsubshs VRT,VRA,VRB

```
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 1856 \\
\hline 1 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
aop}\leftarrow\operatorname{EXTS}((VRA) i:i+15
bop}\leftarrow\operatorname{EXTS}((VRB\mp@subsup{)}{i:i+15}{)
temp }\leftarrow\mathrm{ aop +int }\urcorner\mathrm{ bop +int 1
VRT
end

```

For each integer value i from 0 to 7 , do the following.
Signed-integer halfword element \(i\) in VRB is subtracted from signed-integer halfword element i in VRA.
- If the intermediate result is greater than \(2^{15}-1\) the result saturates to \(2^{15}-1\).
- If the intermediate result is less than \(-2^{15}\) the result saturates to \(-2^{15}\).

The low-order 16 bits of the result are placed into halfword element i of VRT.

\section*{Special Registers Altered:}

SAT

\section*{Version 3.0}

\section*{Vector Subtract Signed Word Saturate} VX-form
```

vsubsws VRT,VRA,VRB

```
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1920} & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
aop}\leftarrow\operatorname{EXTS}((VRA) i:i+31
bop}\leftarrow\operatorname{EXTS}((VRB) (i:i+31
VRT}\mp@subsup{T}{i:i+31}{}\leftarrow\mathrm{ Clamp (aop +int }7\mathrm{ bop +int 1,-2 31, 231}-1
end

```

For each integer value i from 0 to 3 , do the following.
Signed-integer word element i in VRB is subtracted from signed-integer word element \(i\) in VRA.
- If the intermediate result is greater than \(2^{31}-1\) the result saturates to \(2^{31}-1\).
- If the intermediate result is less than \(-2^{31}\) the result saturates to \(-2^{31}\).

The low-order 32 bits of the result are placed into word element i of VRT.

Special Registers Altered: SAT

\section*{Vector Subtract Unsigned Byte Modulo VX-form}

\section*{vsububm \\ VRT,VRA,VRB}
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 1024 \\
\hline 0 & & 6 & & 11 \\
\hline
\end{tabular}
```

do i=0 to 127 by 8
aop}\leftarrow\operatorname{EXTZ}((VRA\mp@subsup{)}{i:i+7}{}
bop }\leftarrow\operatorname{EXTZ((VRB)
VRT i:i+7}<< Chop( aop + int ᄀbop + +int 1, 8),
end

```

For each integer value i from 0 to 15 , do the following. Unsigned-integer byte element \(i\) in VRB is subtracted from unsigned-integer byte element i in VRA. The low-order 8 bits of the result are placed into byte element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Subtract Unsigned Doubleword Modulo VX-form}
```

vsubudm VRT,VRA,VRB

```
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & VRT & VRA & VRB & & 1216 \\
\hline 11 & & 6 & & & \\
\hline
\end{tabular}
```

do i = 0 to 1
aop }\leftarrow\textrm{VR[VRA].dword[i]
bop }\leftarrow\textrm{VR[VRB].dword[i]
VR[VRT].dword[i] \leftarrow Chop( aop +int ~bop + +int 1, 64 )
end

```

For each integer value i from 0 to 1 , do the following. The integer value in doubleword element i of VR[VRB] is subtracted from the integer value in doubleword element \(i\) of VR[ VRA].

The low-order 64 bits of the result are placed into doubleword element \(i\) of VR[VRT].

\section*{Special Registers Altered:}

None

\section*{Programming Note}
vsubudm can be used for signed or unsigned integers.

\section*{Vector Subtract Unsigned Halfword Modulo VX-form}
```

vsubuhm VRT,VRA,VRB

```
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{1088} & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
aop }\leftarrow\operatorname{EXTZ}((VRA\mp@subsup{)}{i:i+15}{}
bop }\leftarrow\operatorname{EXTZ((VRB) i:i+15)
VRT i:i+16}<<Chop( aop + int \negbop + int 1, 16 )
end

```

For each integer value i from 0 to 7 , do the following. Unsigned-integer halfword element \(i\) in VRB is subtracted from unsigned-integer halfword element \(i\) in VRA. The low-order 16 bits of the result are placed into halfword element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Subtract Unsigned Word Modulo VX-form}
vsubuwm VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & VRT & VRA & VRB & & 1152 \\
\hline 11 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
aop \leftarrow EXTZ((VRA) i:i+31)
bop \leftarrow EXTZ((VRB) i:i+31)
VRT}\mp@subsup{T}{i:i+31}{}\leftarrowChop( aop +int fbop +int 1,32
end

```

For each integer value i from 0 to 3 , do the following. Unsigned-integer word element \(i\) in VRB is subtracted from unsigned-integer word element i in VRA. The low-order 32 bits of the result are placed into word element i of VRT.

\section*{Special Registers Altered: \\ None}

\section*{Vector Subtract Unsigned Byte Saturate VX-form \\ vsububs VRT,VRA,VRB}
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1536} & \\
\hline 0 & & 6 & & 11 \\
\hline
\end{tabular}
```

do i=0 to 127 by 8
aop}\leftarrow\operatorname{EXTZ ((VRA)
bop}\leftarrow\operatorname{EXTZ}((VRB\mp@subsup{)}{i:i+7}{\prime}
VRT i:i+7}<<Clamp(aop +int רbop + int 1, 0, 255) 24:31
end

```

For each integer value i from 0 to 15 , do the following.
Unsigned-integer byte element \(i\) in VRB is subtracted from unsigned-integer byte element i in VRA. If the intermediate result is less than 0 the result saturates to 0 . The low-order 8 bits of the result are placed into byte element i of VRT.

\section*{Special Registers Altered:}

SAT

\section*{Vector Subtract Unsigned Halfword Saturate VX-form}
vsubuhs VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1600} & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
aop }\leftarrow\operatorname{EXTZ}((VRA\mp@subsup{)}{i:i+15}{}
bop }\leftarrow\operatorname{EXTZ((VRB) i:i+15)
VRT
end

```

For each integer value i from 0 to 7 , do the following. Unsigned-integer halfword element \(i\) in VRB is subtracted from unsigned-integer halfword element i in VRA. If the intermediate result is less than 0 the result saturates to 0 . The low-order 16 bits of the result are placed into halfword element i of VRT.

\section*{Special Registers Altered:}

SAT

\section*{Vector Subtract Unsigned Word Saturate VX-form}
vsubuws VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{1664} \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
aop}\leftarrow\operatorname{EXTZ ((VRA) i:i+31)
bop }\leftarrow\operatorname{EXTZ ((VRB) i:i+31)
VRT}\mp@subsup{\mp@code{i:i+31}}{}{~Clamp(aop +int }\urcorner\mathrm{ bop + +int 1, 0, 2 22-1)
end

```

For each integer value i from 0 to 7 , do the following. Unsigned-integer word element i in VRB is subtracted from unsigned-integer word element i in VRA.
- If the intermediate result is less than 0 the result saturates to 0 .

The low-order 32 bits of the result are placed into word element i of VRT.

\section*{Special Registers Altered:}

\section*{SAT}

\section*{Vector Subtract Unsigned Quadword Modulo VX-form}
```

vsubuqm VRT,VRA,VRB

```
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1280} \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
src1 }\leftarrow\textrm{VR[VRA]
src2 }\leftarrow\textrm{VR[VRB]
sum \leftarrow EXTZ(src1) + EXTZ(`src2) + EXTZ(1)
VR[VRT] }\leftarrow\mathrm{ Chop(sum, 128)

```

Let srcl be the integer value in VR[ VRA].
Let src2 be the integer value in VR[ VRB].
srcl and srcl can be signed or unsigned integers.
The rightmost 128 bits of the sum of srcl , the one's complement of src , and the value 1 are placed into VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Subtract Extended Unsigned Quadword Modulo VA-form}
vsubeuqm VRT,VRA,VRB,VRC
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & VRC & \multicolumn{2}{|c|}{62} & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
Src1 }\leftarrow\textrm{VR[VRA]
src2 \leftarrow VR[VRB]
cin }\leftarrow\mathrm{ VR[VRC].bit[127]
sum \leftarrow EXTZ(src1) + EXTZ (ᄀsrc2) + EXTZ(cin)
VR[VRT] \& Chop(sum, 128)

```

Let srcl be the integer value in VR[VRA].
Let \(s r_{1}\) be the integer value in VR[VRB].
Let \(\mathrm{c} i n\) be the integer value in bit 127 of VR[ VRC].
\(\operatorname{srcl}\) and \(\operatorname{src}\) can be signed or unsigned integers.

The rightmost 128 bits of the sum of srcl , the one's complement of \(\operatorname{src} 2\), and cin are placed into VR[VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Subtract \& write Carry Unsigned Quadword VX-form}
vsubcuq VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1344} & \\
\hline 0 & & 6 & & & 21 & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
src1 }\leftarrow VR[VRA
src2 }\leftarrow\textrm{VR[VRB]
sum \leftarrowEXTZ(src1) + EXTZ(\negsrc2) + EXTZ(1)
VR[VRT] \leftarrowChop( EXTZ(Chop(sum >> 128, 1) ), 128)

```

Let \(\mathrm{s} r \mathrm{Cl}\) be the integer value in VR[ VRA].
Let \(\operatorname{src} 2\) be the integer value in VR[ VRB].
srcl and srcl can be signed or unsigned integers.
The carry out of the sum of srcl, the one's complement of \(\operatorname{src}\), and the value 1 is placed into VR[VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Subtract Extended \& write Carry Unsigned Quadword VA-form}
vsubecuq VRT,VRA,VRB,VRC
\begin{tabular}{|l|l|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & VRC & & 63 & \\
\hline 0 & & 6 & & 11 & & & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
src1 $\leftarrow \mathrm{VR}$ [VRA]
$\operatorname{src} 2 \leftarrow \mathrm{VR}[\mathrm{VRB}]$
$\operatorname{cin} \leftarrow$ VR[VRC].bit[127]
sum $\leftarrow \operatorname{EXTZ}(\mathrm{src} 1)+\mathrm{EXTZ}(\mathrm{src} 2)+\mathrm{EXTZ}(\mathrm{cin})$
$\operatorname{VR}[V R T] \leftarrow$ Chop ( EXTZ (Chop (sum >> 128, 1) ), 128 )

```

Let srcl be the integer value in VR[ VRA].
Let \(s r_{1}\) be the integer value in VR[ VRB].
Let \(\mathrm{c} i n\) be the integer value in bit 127 of VR[ VRC].
srcl and srcl can be signed or unsigned integers.
The carry out of the sum of srcl, the one's complement of src2, and cin are placed into VR[VRT].

\section*{Special Registers Altered:}

None

\section*{Version 3.0}

\section*{Programming Note}

The Vector Subtract Unsigned Quadword instructions support efficient wide-integer subtraction. The following code sequence can be used to implement a 512-bit signed or unsigned subtract operation.
\begin{tabular}{|c|c|c|}
\hline vsubuqm & vS3, vA3, vB3 & \# bits 384:511 of difference \\
\hline vsubcuq & \(\mathrm{vC} 3, \mathrm{vA} 3, \mathrm{vB} 3\) & \# carry out of bit 384 of difference \\
\hline vsubeuqm & vS2, vA2, vB2, vC3 & \# bits 256:383 of difference \\
\hline vsubecuq & vC2, vA2, vB2, vC3 & \# carry out of bit 256 of difference \\
\hline vsubeuqm & vS1, vA1, vB1, vC2 & \# bits 128:255 of difference \\
\hline vsubecuq & vC1, vA1, vB1, vC2 & \# carry out of bit 128 of difference \\
\hline vsubeuqm & vS0, vA0, vB0, vC1 & \# bits 0:127 of difference \\
\hline
\end{tabular}

\subsection*{6.9.1.3 Vector Integer Multiply Instructions}

\section*{Vector Multiply Even Signed Byte VX-form}
\[
\text { vmulesb } \quad \text { VRT,VRA,VRB }
\]
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 776 \\
\hline 0 & & 61 \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
prod}\leftarrow\operatorname{EXTS}((VRA) i:i+7 ) < < Si EXTS((VRB) i:i+7
VRT i:i+15
end

```

For each integer value i from 0 to 7 , do the following. Signed-integer byte element \(i \times 2\) in VRA is multiplied by signed-integer byte element \(\mathrm{i} \times 2\) in VRB. The low-order 16 bits of the product are placed into halfword element i VRT.

Special Registers Altered:
None

\section*{Vector Multiply Even Unsigned Byte VX-form}
vmuleub VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & VRT & VRA & VRB & & 520 \\
\hline 0 & & 61 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
prod}\leftarrow\operatorname{EXTZ}((VRA) i:i+7 ) < xui EXTZ ((VRB) i:i+7
VRT
end

```

For each integer value i from 0 to 7 , do the following. Unsigned-integer byte element \(\mathrm{i} \times 2\) in VRA is multiplied by unsigned-integer byte element \(\mathrm{i} \times 2\) in VRB. The low-order 16 bits of the product are placed into halfword element i VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Multiply Odd Signed Byte VX-form}
vmulosb VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{264} \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
prod}\leftarrow\operatorname{EXTS}((VRA) (V+8:i+15) < < Si EXTS((VRB) i+8:i+15)
VRT i:i+15
end

```

For each integer value i from 0 to 7 , do the following. Signed-integer byte element \(\mathrm{i} \times 2+1\) in VRA is multiplied by signed-integer byte element \(i \times 2+1\) in VRB. The low-order 16 bits of the product are placed into halfword element i VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Multiply Odd Unsigned Byte VX-form}
vmuloub VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \(\bigcirc\) & \({ }_{6}\) VRT & \({ }_{11}\) VRA & \({ }_{16}\) VRB & 21 & 8 & 31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } \mathrm{i}=0 \text { to } 127 \text { by } 16 \\
& \quad \operatorname{prod} \leftarrow \operatorname{EXTZ}\left(( \operatorname { V R A } ) _ { \mathrm { i } + 8 : i + 1 5 ) } \mathrm { X } _ { \mathrm { ui } } \operatorname { E X T Z } \left((\mathrm{VRB})_{\mathrm{i}+8: i+15)}\right.\right. \\
& \operatorname{VRT}_{\mathrm{i}: \mathrm{i}+15} \leftarrow \text { Chop }(\operatorname{prod}, 16) \\
& \text { end }
\end{aligned}
\]

For each integer value i from 0 to 7 , do the following. Unsigned-integer byte element \(\mathrm{i} \times 2+1\) in VRA is multiplied by unsigned-integer byte element \(\mathrm{i} \times 2+1\) in VRB. The low-order 16 bits of the product are placed into halfword element i VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Multiply Even Signed Halfword VX-form}

\section*{vmulesh VRT,VRA,VRB}
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{840} \\
\hline 0 & & & 11 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
prod}\leftarrow\operatorname{EXTS}((VRA) i:i+15) < < Si EXTS((VRB) i:i+15 )
VRTi:i+31
end

```

For each integer value i from 0 to 3 , do the following. Signed-integer halfword element \(\mathrm{i} \times 2\) in VRA is multiplied by signed-integer halfword element \(i \times 2\) in VRB. The low-order 32 bits of the product are placed into halfword element i VRT.

\section*{Special Registers Altered: None \\ Vector Multiply Even Unsigned Halfword VX-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|l|}{vmuleuh VRT,VRA,VRB} \\
\hline 04 & \[
6
\] & \[
{ }_{11} \text { VRA }
\] & \[
{ }_{16} \text { VRB }
\] & 21 & 584 & 31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 32 \\
& \qquad \text { prod } \leftarrow \operatorname{EXTZ}\left((\mathrm{VRA})_{i: i+15}\right) x_{\text {ui }} \text { EXTZ }\left((\mathrm{VRB})_{i: i+15}\right) \\
& \left.\quad \operatorname{VRT}_{i: i+31} \leftarrow \text { Chop(prod, } 32\right) \\
& \text { end }
\end{aligned}
\]

For each integer value i from 0 to 3 , do the following. Unsigned-integer halfword element i×2 in VRA is multiplied by unsigned-integer halfword element \(\mathrm{i} \times 2\) in VRB. The low-order 32 bits of the product are placed into halfword element i VRT.

\section*{Special Registers Altered: None}

\section*{Vector Multiply Odd Signed Halfword VX-form}
vmulosh VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & VRA & \multicolumn{2}{|c|}{ VRB } & \multicolumn{2}{|c|}{328} \\
\hline 0 & & 6 & & 11 & & 31 \\
\hline
\end{tabular}
```

do i=0 to 127 by 32

```

```

    VRT
    end

```

For each integer value i from 0 to 3 , do the following. Signed-integer halfword element \(i \times 2+1\) in VRA is multiplied by signed-integer halfword element \(\mathrm{i} \times 2+1\) in VRB. The low-order 32 bits of the product are placed into halfword element i VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Multiply Odd Unsigned Halfword VX-form}
vmulouh VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|ccc|}
\hline 4 & VRT & VRA & VRB & & 72 & \\
\hline 0 & & 61 \\
\hline
\end{tabular}
```

do i=0 to 127 by 32

```

```

    VRT
    end

```

For each integer value i from 0 to 3 , do the following. Unsigned-integer halfword element ix2+1 in VRA is multiplied by unsigned-integer halfword element i \(\times 2+1\) in VRB. The low-order 32 bits of the product are placed into halfword element i VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Multiply Even Signed Word VX-form}

\section*{vmulesw VRT,VRA,VRB}
\begin{tabular}{|c|c|c|c|ccc|}
\hline 4 & VRT & VRA & VRB & & 904 & 31 \\
\hline
\end{tabular}
```

do i = 0 to 1
srcl \leftarrowVR[VRA].word[2xi]
src2 \& VR[VRB].word[2xi]
VR[VRT].dword[i] \&srcl x si srcl
end

```

For each integer value i from 0 to 1, do the following. The signed integer in word element \(2 \times i\) of VR[ VRA] is multiplied by the signed integer in word element \(2 \times i\) of VR[VRB].

The 64-bit product is placed into doubleword element \(i\) of VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Multiply Even Unsigned Word VX-form}


For each integer value i from 0 to 1 , do the following.
The unsigned integer in word element \(2 \times i\) of VR[VRA] is multiplied by the unsigned integer in word element \(2 \times i\) of VR[VRB].

The 64-bit product is placed into doubleword element i of VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Multiply Odd Signed Word VX-form}
```

vmulosw VRT,VRA,VRB

```
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & & 392 \\
\hline 11
\end{tabular}
```

do i = 0 to l
srcl}\leftarrowVR[VRA].word[2xi+1
src2 \leftarrowVR[VRB], word[2xi+1]
VR[VRT].dword[i]}\leftarrow\operatorname{srcl}\mp@subsup{x}{si}{src2
end

```

For each integer value i from 0 to 1, do the following.
The signed integer in word element \(2 \times i+1\) of VR[ VRA] is multiplied by the signed integer in word element \(2 \times i+1\) of VR[ VRB].

The 64-bit product is placed into doubleword element i of VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Multiply Odd Unsigned Word VX-form}
vmulouw VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & \({ }_{16}\) VRB & & 136 & \\
\hline 0 & & & & & \\
\hline
\end{tabular}
```

do i = 0 to 1
srcl \leftarrowVR[VRA], word[2xi+1]
srC2 \leftarrowVR[VRB].word[2xi+1]
VR[VRT].dword[i]}\leftarrow\operatorname{srcl xui srcl
end

```

For each integer value i from 0 to 1 , do the following.
The unsigned integer in word element \(2 \times i+1\) of VR[VRA] is multiplied by the unsigned integer in word element \(2 \times i+1\) of VR[ VRB] .

The 64-bit product is placed into doubleword element i of VR[ VRT].

Special Registers Altered:
None

\section*{Version 3.0}

\section*{Vector Multiply Unsigned Word Modulo VX-form}
vmuluwm VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & VRA & VRB & & 137 & 31 \\
\hline
\end{tabular}
```

    do i =0 to 3
        srcl \leftarrowVR[VRA]. word[i]
        src2 \leftarrowVR[VRB].word[i]
        VR[VRT],word[i] & Chop( srcl xui src2, 32)
    end
    ```

The integer in word element \(i\) of VR[ VRA] is multiplied by the integer in word element \(i\) of \(V\) R[ VRB] .

The least-significant 32 bits of the product are placed into word element i of VR[ VRT] .

\section*{Special Registers Altered:} None

\section*{Programming Note}
vmuluwm can be used for unsigned or signed integers.

\subsection*{6.9.1.4 Vector Integer Multiply-Add/Sum Instructions}

\section*{Vector Multiply-High-Add Signed Halfword Saturate VA-form}
vmhaddshs VRT,VRA, VRB,VRC
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 4 & VRT & VRA & VRB & VRC & \multicolumn{2}{|c|}{32} \\
\hline 0 & & 6 & & & \\
\hline 11 & & & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
prod}\leftarrow\operatorname{EXTS}((VRA) i:i+15) < < si EXTS((VRB) i:i+15)
sum}\leftarrow(\operatorname{prod}>>\mp@subsup{>}{\mathrm{ si }}{}15) +\mp@subsup{}{\mathrm{ int }}{}\operatorname{EXTS}((VRC\mp@subsup{)}{i:i+15}{}
VRT}\mp@subsup{T}{i:i+15}{}\leftarrow\mathrm{ Clamp(sum, -2 15, 215}-1) 16:31
end

```

For each vector element i from 0 to 7 , do the following. Signed-integer halfword element \(i\) in VRA is multiplied by signed-integer halfword element in VRB, producing a 32-bit signed-integer product. Bits 0:16 of the product are added to signed-integer halfword element i in VRC.
- If the intermediate result is greater than \(2^{15}-1\) the result saturates to \(2^{15}-1\).
- If the intermediate result is less than \(-2^{15}\) the result saturates to \(-2^{15}\).

The low-order 16 bits of the result are placed into halfword element i of VRT.

Special Registers Altered:
SAT

\section*{Vector Multiply-High-Round-Add Signed Halfword Saturate VA-form}
vmhraddshs VRT,VRA, VRB,VRC
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 4 & VRT & VRA & VRB & VRC & \multicolumn{2}{|c|}{33} \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do $i=0$ to 127 by 16
temp $\leftarrow \operatorname{EXTS}\left((\mathrm{VRC})_{i: i+15}\right)$
prod $\leftarrow \operatorname{EXTS}\left((V R A)_{i: i+15}\right) x_{\text {si }} \operatorname{EXTS}\left((V R B)_{i: i+15}\right)$
sum $\leftarrow\left(\left(\operatorname{prod}+_{\text {int }} 0 x 0000 \_4000\right) \gg_{\text {si }} 15\right)+_{\text {int }}$ temp
$\mathrm{VRT}_{i: i+15} \leftarrow \operatorname{Clamp}\left(\text { sum, }-2^{15}, 2^{15}-1\right)_{16: 31}$
end

```

For each vector element i from 0 to 7, do the following. Signed-integer halfword element \(i\) in VRA is multiplied by signed-integer halfword element i in VRB, producing a 32-bit signed-integer product. The value \(0 \times 0000 \_4000\) is added to the product, producing a 32-bit signed-integer sum. Bits 0:16 of the sum are added to signed-integer halfword element i in VRC.
- If the intermediate result is greater than \(2^{15}-1\) the result saturates to \(2^{15}-1\).
- If the intermediate result is less than \(-2^{15}\) the result saturates to \(-2^{15}\).

The low-order 16 bits of the result are placed into halfword element i of VRT.

\section*{Special Registers Altered:}

SAT

\section*{Vector Multiply-Low-Add Unsigned Halfword Modulo VA-form}
vmladduhm VRT,VRA,VRB,VRC
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 4 & VRT & VRA & VRB & VRC & \multicolumn{2}{|c|}{34} \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
prod}\leftarrowEXTZ((VRA) i:i+15) < <ui EXTZ((VRB) i:i+15
sum}\leftarrowChop(prod,16) +int (VRC) i:i+1
VRT}\mp@subsup{\textrm{i}}{\textrm{i}+1+15}{}\leftarrow\mathrm{ Chop( sum, 16)
end

```

For each integer value i from 0 to 3 , do the following. Unsigned-integer halfword element \(i\) in VRA is multiplied by unsigned-integer halfword element i in VRB, producing a 32-bit unsigned-integer product. The low-order 16 bits of the product are added to unsigned-integer halfword element \(i\) in VRC.

The low-order 16 bits of the sum are placed into halfword element \(i\) of VRT.

\section*{Special Registers Altered:}

\section*{None}

\section*{Programming Note}
vmladduhm can be used for unsigned or signed-integers.

\section*{Vector Multiply-Sum Unsigned Byte Modulo VA-form}
vmsumubm VRT,VRA,VRB,VRC
\begin{tabular}{|c|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & VRC & \multicolumn{2}{|c|}{36} \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
temp}\leftarrow\operatorname{EXTZ}((VRC) i:i+31
do j=0 to 31 by }

```

```

        temp }\leftarrow temp + int prod
    end
    VRT}\mp@subsup{\textrm{i}:1+31}{}{~
    end

```

For each word element in VRT the following operations are performed, in the order shown.
- Each of the four unsigned-integer byte elements contained in the corresponding word element of VRA is multiplied by the corresponding unsigned-integer byte element in VRB, producing an unsigned-integer halfword product.
- The sum of these four unsigned-integer halfword products is added to the unsigned-integer word element in VRC.
- The unsigned-integer word result is placed into the corresponding word element of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Multiply-Sum Mixed Byte Modulo VA-form}
vmsummbm VRT,VRA,VRB,VRC
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & \multicolumn{1}{|c|}{ VRB } & \multicolumn{2}{|c|}{ VRC } & \multicolumn{2}{|c|}{37} & \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
temp}\leftarrow(VRC) i:i+31
do j=0 to 31 by }

```

```

        temp}\leftarrow\mathrm{ temp +int EXTS(prod)
    end
    VRT i:i+31}< \leftarrow temp
    end

```

For each word element in VRT the following operations are performed, in the order shown.
- Each of the four signed-integer byte elements contained in the corresponding word element of VRA is multiplied by the corresponding unsigned-integer byte element in VRB, producing a signed-integer product.
- The sum of these four signed-integer halfword products is added to the signed-integer word element in VRC.
- The signed-integer result is placed into the corresponding word element of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Multiply-Sum Signed Halfword Modulo VA-form}
vmsumshm VRT,VRA,VRB,VRC
\begin{tabular}{|c|c|c|c|c|cc|}
\hline 4 & 6 & VRT & VRA & VRB & VRC & \multicolumn{2}{|c|}{40} & \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do $\mathrm{i}=0$ to 127 by 32
temp $\leftarrow(\mathrm{VRC})_{i: i+31}$
do $j=0$ to 31 by 16
$\operatorname{prod}_{0: 31} \leftarrow(V R A)_{i+j: i+j+15} x_{\text {si }}(\text { VRB })_{i+j: i+j+15}$
temp $\leftarrow$ temp $+_{\text {int }}$ prod
end
$\operatorname{VRT}_{i: i+31} \leftarrow$ temp
end

```

For each word element in VRT the following operations are performed, in the order shown.
- Each of the two signed-integer halfword elements contained in the corresponding word element of VRA is multiplied by the corresponding signed-integer halfword element in VRB, producing a signed-integer product.
- The sum of these two signed-integer word products is added to the signed-integer word element in VRC.
- The signed-integer word result is placed into the corresponding word element of VRT.

Special Registers Altered:
None

\section*{Vector Multiply-Sum Signed Halfword Saturate VA-form}
vmsumshs VRT,VRA,VRB,VRC
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & VRA & VRB & VRC & \multicolumn{2}{|c|}{41} & \\
\hline 0 & & 6 & & 11 & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
temp}\leftarrow\operatorname{EXTS}((VRC) i:i+31
do j=0 to 31 by 16
SrCA}\leftarrow\operatorname{EXTS}((VRA) i+j:i+j+15
srcB}\leftarrow\operatorname{EXTS}((VRB) i+j:i+j+15
prod}\leftarrow\mathrm{ srcA x si srcB
temp }\leftarrow\mathrm{ temp +int prod
end
VRT
end

```

For each word element in VRT the following operations are performed, in the order shown.
- Each of the two signed-integer halfword elements contained in the corresponding word element of VRA is multiplied by the corresponding signed-integer halfword element in VRB, producing a signed-integer product.
- The sum of these two signed-integer word products is added to the signed-integer word element in VRC.
- If the intermediate result is greater than \(2^{31}-1\) the result saturates to \(2^{31}-1\) and if it is less than \(-2^{31}\) it saturates to \(-2^{31}\).
- The result is placed into the corresponding word element of VRT.

\section*{Special Registers Altered:} SAT

\section*{Vector Multiply-Sum Unsigned Halfword Modulo VA-form}
```

vmsumuhm VRT,VRA,VRB,VRC

```
\begin{tabular}{|c|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & VRC & \multicolumn{2}{|c|}{38} \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
temp}\leftarrow\operatorname{EXTZ}((VRC) i:i+31
do j=0 to 31 by 16
srcA}\leftarrow\operatorname{EXTZ}((VRA\mp@subsup{)}{i+j:i+j+15}{}
srcB }\leftarrow\operatorname{EXTZ}((VRB) (i+j:i+j+15
prod}\leftarrow\operatorname{srcA}\mp@subsup{x}{ui}{}\operatorname{srcB
temp }\leftarrow\mathrm{ temp +int prod
end
VRT
end

```

For each word element in VRT the following operations are performed, in the order shown.
- Each of the two unsigned-integer halfword elements contained in the corresponding word element of VRA is multiplied by the corresponding unsigned-integer halfword element in VRB, producing an unsigned-integer word product.
- The sum of these two unsigned-integer word products is added to the unsigned-integer word element in VRC.
- The unsigned-integer result is placed into the corresponding word element of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Multiply-Sum Unsigned Halfword Saturate VA-form}
vmsumuhs VRT,VRA,VRB,VRC
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & VRC & \multicolumn{2}{|c|}{39} \\
\hline 0 & & 6 & & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
temp}\leftarrowEEXTZ ((VRC) i:i+31
do j=0 to 31 by 16
src1 \leftarrow EXTZ((VRA) i+j:i+j+15)
src2 \leftarrow EXTZ((VRB) i+i:i+j+15}
prod }\leftarrow src1 \mp@subsup{x}{ui}{} src
end
temp }\leftarrow\mathrm{ temp + int prod
VRT
end

```

For each word element in VRT the following operations are performed, in the order shown.
- Each of the two unsigned-integer halfword elements contained in the corresponding word element of VRA is multiplied by the corresponding unsigned-integer halfword element in VRB, producing an unsigned-integer product.
- The sum of these two unsigned-integer word products is added to the unsigned-integer word element in VRC.
- If the intermediate result is greater than \(2^{32}-1\) the result saturates to \(2^{32}-1\).
- The result is placed into the corresponding word element of VRT.

Special Registers Altered: SAT

\subsection*{6.9.1.5 Vector Integer Sum-Across Instructions}

\section*{Vector Sum across Signed Word Saturate VX-form}
vsumsws VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1928} & \\
\hline 0 & & 6 & & & & \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { temp } \leftarrow \operatorname{EXTS}\left((\mathrm{VRB})_{96: 127}\right) \\
& \text { do } i=0 \text { to } 127 \text { by } 32 \\
& \quad \text { temp } \leftarrow \text { temp +int EXTS }\left((\mathrm{VRA})_{i: i+31}\right) \\
& \text { end } \\
& \mathrm{VRT}_{0: 31} \leftarrow 0 \times 0000 \_0000 \\
& \mathrm{VRT}_{32: 63} \leftarrow 0 \times 0000 \_0000 \\
& \mathrm{VRT}_{64: 95} \leftarrow 0 x 0000 \_0000 \\
& \left.\mathrm{VRT}_{96: 127} \leftarrow \text { Clamp (temp, }-2^{31}, 2^{31}-1\right)
\end{aligned}
\]

The sum of the four signed-integer word elements in VRA is added to signed-integer word element 3 of VRB.
- If the intermediate result is greater than \(2^{31}-1\) the result saturates to \(2^{31}-1\).
- If the intermediate result is less than \(-2^{31}\) the result saturates to \(-2^{31}\).

The low-end 32 bits of the result are placed into word element 3 of VRT.

Word elements 0 to 2 of VRT are set to 0 .

\section*{Special Registers Altered:} SAT

\section*{Vector Sum across Half Signed Word Saturate VX-form}
vsum2sws VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{1672} & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 64
temp}\leftarrow\operatorname{EXTS}((VRB\mp@subsup{)}{i+32:i+63}{}
do j=0 to 63 by }3
temp}\leftarrow temp +int EXTS ((VRA) i+j:i+j+31
end
VRT}\mp@subsup{\textrm{i}:1+63}{}{\leftarrow}0\times00000_0000|\mathrm{ Clamp(temp, -2 31, 2 31}-1
end

```

Word elements 0 and 2 of VRT are set to 0 .
The sum of the signed-integer word elements 0 and 1 in VRA is added to the signed-integer word element in bits 32:63 of VRB.
- If the intermediate result is greater than \(2^{31}-1\) the result saturates to \(2^{31}-1\).
- If the intermediate result is less than \(-2^{31}\) the result saturates to \(-2^{31}\).

The low-order 32 bits of the result are placed into word element 1 of VRT.

The sum of signed-integer word elements 2 and 3 in VRA is added to the signed-integer word element in bits 96:127 of VRB.
- If the intermediate result is greater than \(2^{31}-1\) the result saturates to \(2^{31}-1\).
- If the intermediate result is less than \(-2^{31}\) the result saturates to \(-2^{31}\).

The low-order 32 bits of the result are placed into word element 3 of VRT.

\section*{Special Registers Altered:} SAT

\section*{Vector Sum across Quarter Signed Byte Saturate VX-form}
vsum4sbs VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 1800 \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
temp}\leftarrow\operatorname{EXTS}((VRB\mp@subsup{)}{i:i+31}{}
do j=0 to 31 by }
temp}\leftarrow\mathrm{ temp +int EXTS((VRA) (i+j:i+j+7)
end
VRT i:i+31}\leftarrow\leftarrow\mathrm{ Clamp(temp, -2 31, 231}-1
end

```

For each integer value i from 0 to 3 , do the following.
The sum of the four signed-integer byte elements contained in word element i of VRA is added to signed-integer word element i in VRB.
- If the intermediate result is greater than \(2^{31}-1\) the result saturates to \(2^{31}-1\).
- If the intermediate result is less than \(-2^{31}\) the result saturates to \(-2^{31}\).

The low-order 32 bits of the result are placed into word element i of VRT.

Special Registers Altered:
SAT

\section*{Vector Sum across Quarter Signed Halfword Saturate VX-form}
```

vsum4shs VRT,VRA,VRB

```
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1608} \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
do \(i=0\) to 127 by 32
temp \(\leftarrow \operatorname{EXTS}\left((\mathrm{VRB})_{i: i+31}\right)\)
do \(j=0\) to 31 by 16
temp \(\leftarrow\) temp \(+_{\text {int }} \operatorname{EXTS}\left((V R A)_{i+j: i+j+15}\right)\)
end
\(\operatorname{VRT}_{i: i+31} \leftarrow\) Clamp (temp, \(\left.-2^{31}, 2^{31}-1\right)\)
end
For each integer value i from 0 to 3 , do the following.
The sum of the two signed-integer halfword elements contained in word element \(i\) of VRA is added to signed-integer word element i in VRB.
- If the intermediate result is greater than \(2^{31}-1\) the result saturates to \(2^{31}-1\).
- If the intermediate result is less than \(-2^{31}\) the result saturates to \(-2^{31}\).

The low-order 32 bits of the result are placed into the corresponding word element of VRT.

\section*{Special Registers Altered:}

SAT

\section*{Version 3.0}

\section*{Vector Sum across Quarter Unsigned} Byte Saturate VX-form
vsum4ubs VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{1544} \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
temp}\leftarrowE\operatorname{EXTZ}((VRB\mp@subsup{)}{i:i+31}{}
do j=0 to 31 by }

```

```

    end
    VRTi:i+31
    end

```

For each integer value i from 0 to 3 , do the following.
The sum of the four unsigned-integer byte elements contained in word element \(i\) of VRA is added to unsigned-integer word element i in VRB.
- If the intermediate result is greater than \(2^{32}-1\) it saturates to \(2^{32}-1\).

The low-order 32 bits of the result are placed into word element \(i\) of VRT.

Special Registers Altered:
SAT

\subsection*{6.9.1.6 Vector Integer Negate Instructions}

\section*{Vector Negate Word VX-form}
Vnegw
\begin{tabular}{|l|l|l|l|l|lll|}
\hline 4 & VRT,VRB \\
0 & & 6 & 6 & \multicolumn{2}{|c|}{ VRB } & & 1538 \\
\hline 11 & & 16 & 21 & & 31 \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailable()
do \(i=0\) to 3
src \(\leftarrow\) EXTS (VR[VRB], word[i])
VR[VRT] word[i] \(\leftarrow\) Chopl \((\operatorname{ssc} \mathrm{c}+1)\), 32\()\)
end
For each integer value i from 0 to 3 , do the following.
The sum of the one's-complement of the signed integer in word element i of VR[VRB] and 1 is placed into word element \(i\) of VR[ VRT].

Special Registers Altered:
None

\section*{Vector Negate Doubleword VX-form}
vnegd VRT,VRB
\begin{tabular}{|l|l|ll|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & \multicolumn{2}{|c|}{7} & \multicolumn{2}{|c|}{ VRB } & \\
31 & 1538 & \\
\hline
\end{tabular}
if MSR.VEC=O then Vector Unavailable(l)
do \(i=0\) to 1
sic \(\leftarrow\) EXTS (VR[VRB]. dword[i])
VR[VRT]dword[i] \(\leftarrow\) Chop( \((\operatorname{sic} \mathrm{C}+1), 64)\)
end
For each integer value i from 0 to 1 , do the following.
The sum of the one's-complement of the signed integer in doubleword element \(i\) of VR[VRB] and 1 is placed into doubleword element \(i\) of VR[ VRT].

Special Registers Altered:
None

\subsection*{6.9.2 Vector Extend Sign Instructions}
\begin{tabular}{l} 
Vector Extend Sign Byte To Word VX-form \\
vextsb2w VRT,VRB \\
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 4 & VRT & 16 & VRB & & 1538 & \\
\hline 0 & & 6 & 11 & & 16 & 21
\end{tabular} \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailablel)
do \(\mathrm{i}=0\) to 3
VR[VRT], word[i] EXTS32(VR[VRB], word[i], byte[3])
end
For each integer value i from 0 to 3 , do the following. The rightmost byte of word element \(i\) of VR[VRB] is sign-extended and placed into word element \(i\) of VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Extend Sign Halfword To Word VX-form \\ vextsh2w VRT,VRB}
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & 17 & VRB & & 1538 & \\
\hline 0 & & 61 & & & \\
\hline
\end{tabular}
if MSR. VEC=O then Vector_Unavailablel)
do \(i=0\) to 3
VR[ VRT], word[i] EXTS32 [VR[VBB]. wor d[i]. hwor d[ 1])
end
For each integer value i from 0 to 3 , do the following.
The rightmost halfword of word element i of VR[VRB] is sign-extended and placed into word element i of VR[VRT].

Special Registers Altered:
None

\section*{Vector Extend Sign Byte To Doubleword VX-form}
vextsb2d
\begin{tabular}{|l|l|l|l|l|lll|}
\hline 4 & VRT,VRB \\
0 & & 6 & 24 & VRB & & 1538 & \\
\hline 11 & & 16 & 21 & & 31 \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable(l
do i = 0 to 1
VR[VRT].dword[i] \& EXTS64(VR[VRB].dword[i]. byte[7])
end

```

For each integer value i from 0 to 1 , do the following. The rightmost byte of doubleword element i of VR[VRB] is sign-extended and placed into doubleword element \(i\) of VR[ VRT].

Special Registers Altered: None

\section*{Vector Extend Sign Halfword To Doubleword VX-form}
vextsh2d VRT,VRB
\begin{tabular}{|l|l|l|l|l|ll|}
\hline 4 & VRT & \multicolumn{1}{|c|}{25} & VRB & & 1538 & \\
\hline 0 & & 6 & & 16 & 21 & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable|)
if "vextsh2d" then do i = 0 to 1
VR[VRT].dword[i] \& ExT564[VR[ VRB], dwor d[i], hwor d[ 3])
end

```

For each integer value i from 0 to 1 , do the following. The rightmost halfword of doubleword element i of VR[VRB] is sign-extended and placed into doubleword element \(i\) of VR[VRT].

Special Registers Altered:
None

Vector Extend Sign Word To Doubleword VX-form
vextsw2d VRT,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & 26 & VRB & & 1538 & \\
\hline 0 & & 61 & & & \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailable()
do \(i=0\) to 1
VR[ VRT]. dword[i] e EXTS64 [VR[ VRB]. dwor d[i]. wor d[1])
end
For each integer value i from 0 to 1, do the following. The rightmost word of doubleword element i of VR[VRB] is sign-extended and placed into doubleword element \(i\) of VR[VRT].

Special Registers Altered:
None

\subsection*{6.9.2.1 Vector Integer Average Instructions}

\section*{Vector Average Signed Byte VX-form}
vavgsb VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{1282} \\
\hline 0 & & 61 \\
\hline
\end{tabular}
```

do i=0 to 127 by 8
aop }\leftarrow\operatorname{EXTS}((VRA) i:i+7 )
bop}\leftarrow\operatorname{EXTS}((VRB) (i:i+7
VRT}\mp@subsup{T}{i:i+7}{}\leftarrow\mathrm{ Chop(( aop +int bop +int 1)>> 1, 8)
end

```

For each integer value i from 0 to 15 , do the following.
Signed-integer byte element \(i\) in VRA is added to signed-integer byte element \(i\) in VRB. The sum is incremented by 1 and then shifted right 1 bit.

The low-order 8 bits of the result are placed into byte element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Average Signed Halfword VX-form}
vavgsh VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{1346} \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 16 \\
& \quad \text { aop } \leftarrow \operatorname{EXTS}\left((\operatorname{VRA})_{i: i+15}\right) \\
& \quad \text { bop } \leftarrow \operatorname{EXTS}\left((\operatorname{VRB})_{i: i+15}\right) \\
& \operatorname{VRT}_{i: i+15} \leftarrow \operatorname{Chop}\left(\left(\text { aop }+_{\text {int }} \text { bop }+_{\text {int }} 1\right) \gg 1,16\right) \\
& \text { end }
\end{aligned}
\]

For each integer value i from 0 to 7 , do the following. Signed-integer halfword element \(i\) in VRA is added to signed-integer halfword element \(i\) in VRB. The sum is incremented by 1 and then shifted right 1 bit.

The low-order 16 bits of the result are placed into halfword element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Average Signed Word VX-form}
vavgsw VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|ccc|}
\hline 4 & VRT & VRA & VRB & & 1410 & \\
\hline 0 & & 6 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by }3
aop}\leftarrow\operatorname{EXTS}((VRA) i:i+31
bop }\leftarrow\operatorname{EXTS}((VRB\mp@subsup{)}{i:i+31}{}
VRT
end

```

For each integer value i from 0 to 3 , do the following.
Signed-integer word element \(i\) in VRA is added to signed-integer word element \(i\) in VRB. The sum is incremented by 1 and then shifted right 1 bit.

The low-order 32 bits of the result are placed into word element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Average Unsigned Byte VX-form}
vavgub VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{1026} \\
\hline 0 & & 61 \\
\hline
\end{tabular}
```

do i=0 to 127 by }
aop}\leftarrow\operatorname{EXTZ}((\textrm{VRA}\mp@subsup{)}{:i+7}{}
bop}\leftarrow\operatorname{EXTZ}((VRB) i:i+7
VRT i:i+7}<<<Chop((aop + int bop +int 1) >> wi 1, 8
end

```

For each integer value i from 0 to 15 , do the following.
Unsigned-integer byte element \(i\) in VRA is added to unsigned-integer byte element i in VRB. The sum is incremented by 1 and then shifted right 1 bit.

The low-order 8 bits of the result are placed into byte element i of VRT.

Special Registers Altered:
None

\section*{Vector Average Unsigned Word VX-form}
```

vavguw VRT,VRA,VRB

```
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & & 1154 & \\
\hline 0 & & 6 & 11 & & 21 & \\
\hline
\end{tabular}

For each integer value i from 0 to 3 , do the following. Unsigned-integer word element \(i\) in VRA is added to unsigned-integer word element i in VRB. The sum is incremented by 1 and then shifted right 1 bit.

The low-order 32 bits of the result are placed into word element i of VRT.

Special Registers Altered:
None
```

do i=0 to 127 by 32

```
do i=0 to 127 by 32
    aop }\leftarrow\operatorname{EXTZ}((VRA\mp@subsup{)}{i:i+31}{}
    aop }\leftarrow\operatorname{EXTZ}((VRA\mp@subsup{)}{i:i+31}{}
    bop }\leftarrow\operatorname{EXTZ ((VRB) i:i+31}
    bop }\leftarrow\operatorname{EXTZ ((VRB) i:i+31}
    VRT i:i+31}\leftarrowCChop((aop + int bop +int 1) >> ui 1, 32)
    VRT i:i+31}\leftarrowCChop((aop + int bop +int 1) >> ui 1, 32)
end
```

end

```

\section*{Vector Average Unsigned Halfiword VX-form}

\section*{vavguh VRT,VRA,VRB}
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1090} \\
\hline 0 & & & & 11 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
aop}\leftarrow\operatorname{EXTZ}((VRA) i:i+15
bop }\leftarrow\operatorname{EXTZ((VRB)
VRT i:i+15
end

```

For each integer value i from 0 to 7 , do the following. Unsigned-integer halfword element \(i\) in VRA is added to unsigned-integer halfword element \(i\) in VRB. The sum is incremented by 1 and then shifted right 1 bit.

The low-order 16 bits of the result are placed into halfword element i of VRT.

\section*{Special Registers Altered: \\ None}

\subsection*{6.9.2.2 Vector Integer Absolute Difference Instructions}

This section describes a set of instructions that return the absolute value of the difference of integer values.

\section*{Vector Absolute Difference Unsigned Byte VX-form}
Vabsdub
\begin{tabular}{|c|c|c|c|cc|}
\hline 4 & VRT, VRA, VRB \\
\hline 0 & & 6 & VRA & VRB & \\
\hline 11 & 1027 & & \\
\hline
\end{tabular}
```

if MSR.VEC=O then Vector_Unavailable()
for i = 0 to 15
srcl \leftarrow EXTZ(VR[VRA],byte[i])
srcl \leftarrowEXTZ(VR[VRB],byte[i])
if (srcl>src2) then
VR[VRT],byte[i] <Chop(srcl + -srcl +1,8)
else
VR[ VRT], byte[i] \leftarrowChop(srcl + -5rcl +1, 8)
end

```

For each integer value i from 0 to 15, do the following. The unsigned integer value in byte element i of VR[VRA] is subtracted by the unsigned integer value in byte element i of VR[VRB]. The absolute value of the difference is placed into byte element i of VR[VRT].

Special Registers Altered: None

\section*{Vector Absolute Difference Unsigned Halfword VX-form}
vabsduh VRT,VRA,VRB

```

if MSR.VEC=O then Vector_Unavailable()
for i = 0 to 7
srcl \leftarrow EXTZ(VR[VRA].hword[i])
srcl \leftarrow EXTZ[VR[VRB].hword[i])
if (srcl>src2) then
VR[VRT],hword[i] \leftarrowChop(srcl + -srcl + 1, 16)
else
VR[VRT],hword[i] \leftarrowChop(sicl + -sicl + 1, 16)
end

```

For each integer value i from 0 to 7 , do the following. The unsigned integer value in halfword element i of VR[VRA] is subtracted by the unsigned integer value in halfword element i of VR[VRB]. The absolute value of the difference is placed into halfword element \(i\) of VR[VRT].

Special Registers Altered:
None

Vector Absolute Difference Unsigned Word VX-form
vabsduw VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & & 1155 & \\
\hline 0 & & & 11 & & & \\
\hline
\end{tabular}
if MSR.VEC=O then Vector _Unavailable()
for \(\mathrm{i}=0\) to 3
srcl \(\leftarrow\) EXTZ(VR[VRA]. word[i])
\(\operatorname{srcz} \leftarrow\) EXTZ (VR[VRB]. word[i])
if (srcl>src2) then
VR[VRT]. word[i] \(\leftarrow\) Chop(srcl + -srcl \(+1,32)\)
else
VR[VRT].word[i] \(\leftarrow\) Chop(srcl \(+\operatorname{srcl}+1,32)\)
end
For each integer value i from 0 to 3 , do the following. The unsigned integer value in word element \(i\) of VR[VRA] is subtracted by the unsigned integer value in word element \(i\) of VR[ VRB]. The absolute value of the difference is placed into word element i of VR[VRT].

Special Registers Altered: None

\subsection*{6.9.2.3 Vector Integer Maximum and Minimum Instructions}

\section*{Vector Maximum Signed Byte VX-form}

```

do i=0 to 127 by 8
aop}\leftarrow\operatorname{EXTS}((VRA\mp@subsup{)}{i:i+7}{}
bop }\leftarrow\operatorname{EXTS}((VRB\mp@subsup{)}{i:i+7}{}
VRT}\mp@subsup{T}{i:i+7}{}\leftarrow(\mathrm{ aop >
end

```

For each integer value i from 0 to 15 , do the following.
Signed-integer byte element in VRA is compared to signed-integer byte element i in VRB. The larger of the two values is placed into byte element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Maximum Signed Doubleword VX-form}
vmaxsd VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & & 450 & \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

do i = 0 to 1
aop \leftarrowVR[VRA].dword[i]
bop }\leftarrow\textrm{VR[VRB].dword[i]
VR[VRT].dword[i] \leftarrow(aop > si bop) ? aop : bop
end

```

For each integer value i from 0 to 1 , do the following. The signed integer value in doubleword element \(i\) of VR[VRA] is compared to the signed integer value in doubleword element \(i\) of VR[VRB]. The larger of the two values is placed into doubleword element i of VR[VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Maximum Unsigned Byte VX-form}
vmaxub VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & 2 & 2 \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by }
aop}\leftarrow\operatorname{EXTZ ((VRA) i:i+7}
bop }\leftarrow\operatorname{EXTZ}((VRB\mp@subsup{)}{i:i+7}{}
VRT i:i+7}<\leftarrow(\mathrm{ aop >ui bop) ? (VRA) i:i+7 : (VRB) i:i+7
end

```

For each integer value i from 0 to 15 , do the following.
Unsigned-integer byte element \(i\) in VRA is compared to unsigned-integer byte element i in VRB. The larger of the two values is placed into byte element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Maximum Unsigned Doubleword VX-form}
vmaxud VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & & 194 & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i = 0 to 1
aop \leftarrow VR[VRA].dword[i]
bop \leftarrowVR[VRB].dword[i]
VR[VRT].dword[i] \leftarrow(aop >ui bop) ? aop : bop
end

```

For each integer value i from 0 to 1 , do the following. The unsigned integer value in doubleword element \(i\) of VR[VRA] is compared to the unsigned integer value in doubleword element \(i\) of VR[ VRB]. The larger of the two values is placed into doubleword element \(i\) of VR[ VRT].

Special Registers Altered:
None

\section*{Vector Maximum Signed Halfword VX-form}
```

vmaxsh VRT,VRA,VRB

```
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 322 \\
\hline 11 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
aop }\leftarrow\operatorname{EXTS ((VRA) i:i+15)
bop}\leftarrow\operatorname{EXTS (VRB) i:i+15
VRT i:i+15}<<(aop > si bop ) ? (VRA) i:i+15 : (VRB) i:i+15
end

```

For each integer value i from 0 to 7 , do the following. Signed-integer halfword element \(i\) in VRA is compared to signed-integer halfword element i in VRB. The larger of the two values is placed into halfword element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Maximum Signed Word VX-form}
vmaxsw VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{386} & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
aop }\leftarrow\operatorname{EXTS}((VRA) i:i+31
bop }\leftarrow\operatorname{EXTS((VRB) i:i+31)
VRT}\mp@subsup{i}{i:i+31}{}\leftarrow(\mathrm{ aop > si mop ) ? (VRA) (i:i+31 :(VRB) i:i+31
end

```

For each integer value i from 0 to 3, do the following. Signed-integer word element \(i\) in VRA is compared to signed-integer word element \(i\) in VRB. The larger of the two values is placed into word element i of VRT.

Special Registers Altered:
None

\section*{Vector Maximum Unsigned Halfword VX-form}

\section*{vmaxuh VRT,VRA,VRB}
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{66} \\
\hline
\end{tabular}
```

do $i=0$ to 127 by 16
aop $\leftarrow \operatorname{EXTZ}\left((\mathrm{VRA})_{i: i+15}\right)$
bop $\leftarrow \operatorname{EXTZ}\left((V R B)_{i: i+15}\right)$
$V R T_{i: i+15} \leftarrow\left(\right.$ aop $>_{\text {ui }}$ bop) ? $(V R A)_{i: i+15}:(V R B)_{i: i+15}$
end

```

For each integer value i from 0 to 7 , do the following. Unsigned-integer halfword element \(i\) in VRA is compared to unsigned-integer halfword element i in VRB. The larger of the two values is placed into halfword element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Maximum Unsigned Word VX-form}
vmaxuw VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 130 \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
aop }\leftarrow\operatorname{EXTZ}((VRA\mp@subsup{)}{i:i+31}{}
bop }\leftarrow\operatorname{EXTZ ((VRB) i:i+31)
VRT i:i+31}<\leftarrow(\mathrm{ aop >ui bop) ? (VRA)
end

```

For each integer value i from 0 to 3 , do the following. Unsigned-integer word element i in VRA is compared to unsigned-integer word element i in VRB. The larger of the two values is placed into word element i of VRT.

Special Registers Altered:
None

\section*{Vector Minimum Signed Byte VX-form}
Vminsb
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT, VRA, VRB \\
\hline 0 & & 6 & VRA & VRB & \\
\hline 11
\end{tabular}
```

do i=0 to 127 by 8
aop}\leftarrow\operatorname{EXTS}((VRA) i:i+7
bop}\leftarrow\operatorname{EXTS}((VRB\mp@subsup{)}{i:i+7}{\prime}
VRT}\mp@subsup{\mp@code{i:i+7}}{}{\leftarrow(\mathrm{ (aop < si bop) ? (VRA) i:i+7 : (VRB) i:i+7}
end

```

For each integer value i from 0 to 15 , do the following.
Signed-integer byte element in VRA is compared to signed-integer byte element i in VRB. The smaller of the two values is placed into byte element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Minimum Signed Doubleword VX-form}
Vminsd
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT, VRA, VRB \\
\hline 0 & & 6 & VRT & VRA & VRB \\
16 & & 962 & & \\
\hline
\end{tabular}
```

do i = 0 to 1
aop \leftarrowVR[VRA].dword[i]
bop }\leftarrow\operatorname{VR[VRB].dword[i]
VR[VRT].dword[i] \leftarrow(EXTS(aop) < <si EXTS (bop) ) ? aop : bop
end

```

For each integer value i from 0 to 1 , do the following. The signed integer value in doubleword element i of VR[VRA] is compared to the signed integer value in doubleword element i of VR[ VRB]. The smaller of the two values is placed into doubleword element \(i\) of VR[VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Minimum Unsigned Byte VX-form}
Vminub
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT, VRA, VRB \\
0 & & 6 & VRT & VRA & VRB \\
16 & & 514 & \\
\hline
\end{tabular}
```

do i=0 to 127 by 8
aop}\leftarrow\operatorname{EXTZ}((VRA) i:i+7
bop }\leftarrow\operatorname{EXTZ}((VRB\mp@subsup{)}{i:i+7}{
VRT}\mp@subsup{T}{i:i+7}{}\leftarrow(\mathrm{ aop <ui bop ) ?(VRA) i:i+7 : (VRB) i:i+7
end

```

For each integer value i from 0 to 15 , do the following.
Unsigned-integer byte element \(i\) in VRA is compared to unsigned-integer byte element i in VRB. The smaller of the two values is placed into byte element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Minimum Unsigned Doubleword VX-form}
vminud VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & & 706 \\
\hline 0 & & & 11 & & 21 & \\
\hline
\end{tabular}
```

do i = 0 to 1
aop }\leftarrow\mathrm{ VR[VRA].dword[i]
bop }\leftarrowVR[VRB].dword[i]
VR[VRT].dword[i] \& (aop <ui bop) ? aop : bop
end

```

For each integer value i from 0 to 1 , do the following. The unsigned integer value in doubleword element \(i\) of VR[VRA] is compared to the unsigned integer value in doubleword element \(i\) of VR[ VRB]. The smaller of the two values is placed into doubleword element \(i\) of VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Minimum Signed Halfword VX-form}

\section*{vminsh}

VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 834 \\
\hline 11 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
aop}\leftarrow\operatorname{EXTS}((\mathrm{ VRA )
bop }\leftarrow\operatorname{EXTS}((VRB) (:i+15
VRT i:i+15}<\leftarrow(aop<si bop ) ? (VRA) i:i+15 : (VRB) i:i+15
end

```

For each integer value i from 0 to 7 , do the following. Signed-integer halfword element \(i\) in VRA is compared to signed-integer halfword element i in VRB. The smaller of the two values is placed into halfword element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Minimum Signed Word VX-form}
vminsw VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|l|ll|}
\hline 4 & VRT & VRA & VRB & & 898 & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
aop}\leftarrow\operatorname{EXTS}((VRA) i:i+31
bop }\leftarrow\operatorname{EXTS((VRB) (:i+31)
VRT}\mp@subsup{T}{i:i+31}{}\leftarrow(\mathrm{ aop <si mop ) ? (VRA) i:i+31 :(VRB) i:i+31
end

```

For each integer value i from 0 to 3, do the following. Signed-integer word element \(i\) in VRA is compared to signed-integer word element i in VRB. The smaller of the two values is placed into word element i of VRT.

Special Registers Altered:
None

\section*{Vector Minimum Unsigned Halfword VX-form}

\section*{vminuh VRT,VRA,VRB}
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{578} \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do $i=0$ to 127 by 16
aop $\leftarrow \operatorname{EXTZ}\left((V R A)_{i: i+15}\right)$
bop $\leftarrow \operatorname{EXTZ}\left((V R B)_{i: i+15}\right)$
$\operatorname{VRT}_{i: i+15} \leftarrow($ aop <ui bop $)$ ? $(V R A)_{i: i+15}:(V R B)_{i: i+15}$
end

```

For each integer value i from 0 to 7 , do the following. Unsigned-integer halfword element \(i\) in VRA is compared to unsigned-integer halfword element i in VRB. The smaller of the two values is placed into halfword element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Minimum Unsigned Word VX-form}
vminuw VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 642 \\
\hline 10 & & 61 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
aop }\leftarrow\operatorname{EXTZ}((VRA\mp@subsup{)}{i:i+31}{}
bop }\leftarrow\operatorname{EXTZ((VRB)
VRTi:i+31
end

```

For each integer value i from 0 to 3 , do the following. Unsigned-integer word element i in VRA is compared to unsigned-integer word element i in VRB. The smaller of the two values is placed into word element i of VRT.

Special Registers Altered:
None

\subsection*{6.9.3 Vector Integer Compare Instructions}

The Vector Integer Compare instructions compare two Vector Registers element by element, interpreting the elements as unsigned or signed-integers depending on the instruction, and set the corresponding element of the target Vector Register to all 1s if the relation being tested is true and to all 0 s if the relation being tested is false.

If Rc=1 CR Field 6 is set to reflect the result of the comparison, as follows.

\section*{Bit Description}

0 The relation is true for all element pairs (i.e., VRT is set to all 1s)

10
2 The relation is false for all element pairs (i.e., VRT is set to all Os)

30

\section*{Programming Note}
vcmpequb[.], vcmpequh[.], vcmpequw[.], and vcmpequd[.] can be used for unsigned or signed-integers.

\section*{Vector Compare Equal Unsigned Byte VC-form}
\begin{tabular}{lll} 
vcmpequb & VRT,VRA,VRB & \((R c=0)\) \\
vcmpequb. & VRT,VRA,VRB & \((R c=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 4 & VRT & VRA & VRB & Rc & 6 & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } \mathrm{i}=0 \text { to } 127 \text { by } 8 \\
& \mathrm{VRT}_{\mathrm{i}: 1+7} \leftarrow\left((\mathrm{VRA})_{i: i+7}={ }_{\text {int }}(\mathrm{VRB})_{i: i+7}\right) ?{ }^{8} 1:{ }^{8} 0
\end{aligned}
\]

For each integer value i from 0 to 15 , do the following. Unsigned-integer byte element \(i\) in VRA is compared to unsigned-integer byte element i in VRB. Byte element \(i\) in VRT is set to all 1s if unsigned-integer byte element \(i\) in VRA is equal to unsigned-integer byte element \(i\) in VRB, and is set to all Os otherwise.

\section*{Special Registers Altered:}

CR field 6
.(if Rc=1)

\section*{Vector Compare Equal Unsigned Halfword VC-form}
\begin{tabular}{lll} 
vcmpequh & VRT,VRA,VRB & \((\mathrm{Rc}=0)\) \\
vcmpequh. & VRT,VRA,VRB & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & Rc & 70 & \\
\hline 0 & 6 & 11 & 16 & 2122 & & 31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } \mathrm{i}=0 \text { to } 127 \text { by } 16 \\
& \operatorname{VRT}_{i: i+15} \leftarrow\left((\mathrm{VRA})_{i: i+15}=\text { int }(\mathrm{VRB})_{i: i+15}\right) ?{ }^{16} 1:{ }^{16} 0 \\
& \text { end } \\
& \text { if Rc=1 then do } \\
& \mathrm{t} \leftarrow\left(\mathrm{VRT}={ }^{128} 1\right) \\
& \mathrm{f} \leftarrow\left(\mathrm{VRT}=^{128} 0\right) \\
& \text { CR6 } \leftarrow \mathrm{t}||\mathrm{ObO}|| \mathrm{f}|\mid 0 \mathrm{bo}
\end{aligned}
\]

For each integer value i from 0 to 7 , do the following. Unsigned-integer halfword element \(i\) in VRA is compared to unsigned-integer halfword element element in VRB. Halfword element \(i\) in VRT is set to all 1s if unsigned-integer halfword element \(i\) in VRA is equal to unsigned-integer halfword element \(i\) in VRB, and is set to all 0 s otherwise.

\section*{Special Registers Altered:}

CR field 6
(if Rc=1)

\section*{Vector Compare Equal Unsigned Word VC-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{vcmpequw vcmpequw.}} & \multicolumn{3}{|l|}{\multirow[t]{2}{*}{\begin{tabular}{l}
VRT,VRA,VRB \\
VRT,VRA,VRB
\end{tabular}}} & & ( \(R C=0\) ) \\
\hline & & & & & & ( \(\mathrm{RC}=1\) ) \\
\hline 4 & \[
{ }_{6} \text { VRT }
\] & \[
{ }_{11} \text { VRA }
\] & \[
{ }_{16} \text { VRB }
\] & \(\left|\begin{array}{l}\mathrm{Rc} \\ 21\end{array}\right|_{22}\) & 134 & 31 \\
\hline
\end{tabular}
```

do $\mathrm{i}=0$ to 127 by 32
$\operatorname{VRT}_{i: i+31} \leftarrow\left((\mathrm{VRA})_{i: i+31}=_{\text {int }}(\mathrm{VRB})_{i: i+31}\right) ?{ }^{32} 1:{ }^{32} 0$
end
if $\mathrm{Rc}=1$ then do
$t \leftarrow\left(\right.$ VRT $\left.={ }^{128} 1\right)$
$\mathrm{f} \leftarrow\left(\mathrm{VRT}={ }^{128} 0\right)$
$\mathrm{CR6} \leftarrow \mathrm{t}\|\mathrm{ObO}\| \mathrm{f} \| \mathrm{ObO}$
end

```

For each integer value i from 0 to 3 , do the following. The unsigned integer value in word element \(i\) in VR[ VRA] is compared to the unsigned integer value in word element \(i\) in VR[VRB]. Word element \(i\) in VR[ VRT] is set to all 1s if unsigned-integer word element \(i\) in VR[VRA] is equal to unsigned-integer word element i in VR[VRB], and is set to all Os otherwise.

\section*{Special Registers Altered:}

CR field 6 (if \(R C=1\) )

\section*{Vector Compare Equal Unsigned Doubleword VX-form}

```

do i = 0 to 1
aop \leftarrow EXTZ (VR[VRA].dword[i])
bop \leftarrow EXTZ (VR[VRB].dword[i])
if (aop = bop) then do
VR[VRT].dword[i] \leftarrow OxFFFF_FFFF_FFFF_FFFF
flag.bit[i] }\leftarrow0\mathrm{ Ob1
end
else do
VR[VRT].dword[i] \leftarrow 0x0000_0000_0000_0000
flag.bit[i] \leftarrow 0b0
end
end
if Rc=1 then do
CR.bit[24]}\leftarrow(f1ag=0b11
CR.bit[25] }\leftarrow0\mathrm{ 0b0
CR.bit[26] \leftarrow(flag=0b00)
CR.bit[27] \leftarrow0b0
end

```

For each integer value \(i\) from 0 to 1 , do the following.
The unsigned integer value in doubleword element \(i\) of VR[VRA] is compared to the unsigned integer value in doubleword element \(i\) of VR[VRB]. Doubleword element \(i\) of VR[VRT] is set to all 1 s if the unsigned integer value in doubleword element i of VR[VRA] is equal to the unsigned integer value in doubleword element \(i\) of VR[VRB], and is set to all Os otherwise.

Special Registers Altered:
CR field 6
(if \(R C=1\) )

\section*{Vector Compare Greater Than Signed Byte VC-form}
\begin{tabular}{lll} 
vcmpgtsb & VRT,VRA,VRB & \((R C=0)\) \\
vcmpgtsb. & VRT,VRA,VRB & \((R c=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & Rc & 774 & \\
\hline 0 & & 6 & 11 & & & & 22
\end{tabular}
\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 8 \\
& \quad \operatorname{VRT}_{i: i+7} \leftarrow\left((\mathrm{VRA})_{i: i+7}>_{\text {si }}(\mathrm{VRB})_{i: i+7}\right) \quad{ }^{8} 1:{ }^{8} 0 \\
& \text { end } \\
& \text { if RC=1 then do } \\
& \quad t \leftarrow\left(\mathrm{VRT}=1{ }^{128} 1\right) \\
& \quad \mathrm{f} \leftarrow\left(\mathrm{VRT}={ }^{128} 0\right) \\
& \quad \text { CR6 } \leftarrow t\|0 b 0\| f \| 0 \mathrm{O} 0 \\
& \text { end }
\end{aligned}
\]

For each integer value \(i\) from 0 to 15 , do the following.
The signed integer value in byte element i in VR[ VRA] is compared to the signed integer value in byte element \(i\) in VR[VRB]. Byte element \(i\) in VR[VRT] is set to all 1 s if signed-integer byte element i in VR[VRA] is greater than to signed-integer byte element i in VR[ VRB], and is set to all 0 s otherwise.

\section*{Special Registers Altered:}

CR field 6
(if \(R C=1\) )

\section*{Vector Compare Greater Than Signed Doubleword VX-form}
\begin{tabular}{lll} 
vcmpgtsd & VRT,VRA,VRB & \((R c=0)\) \\
vcmpgtsd. & VRT,VRA,VRB & \((R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & \begin{tabular}{l} 
RC \\
21
\end{tabular} & 967 & \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i = 0 to 1
aop \leftarrow EXTS (VR[VRA] .dword[i])
bop \leftarrow EXTS (VR[VRB] .dword[i])
if (aop > si bop) then do
VR[VRT].dword[i] \leftarrow0xFFFF_FFFF_FFFF_FFFF
flag.bit[i] }\leftarrow0\mathrm{ Ob1
end
else do
VR[VRT].dword[i] \leftarrow0x0000_0000_0000_0000
flag.bit[i] }\leftarrow\mathrm{ ObO
end
end
if "vcmpgtsd." then do
CR.bit[24] \leftarrow(flag=0b11)
CR.bit[25] \leftarrow 0b0
CR.bit[26] \leftarrow(flag=0b00)
CR.bit[27] }\leftarrow0\mathrm{ ObO
end

```

For each integer value i from 0 to 1 , do the following.
The signed integer value in doubleword element \(i\) of VR[VRA] is compared to the signed integer value in doubleword element i of VR[ VRB]. Doubleword element \(i\) of VR[ VRT] is set to all 1 s if the signed integer value in doubleword element \(i\) of VR[ VRA] is greater than the signed integer value in doubleword element i of VR[VRB], and is set to all Os otherwise.

Special Registers Altered:
CR field 6
(if \(R C=1\) )

\section*{Vector Compare Greater Than Signed Halfword VC-form}
\begin{tabular}{lll} 
vcmpgtsh & VRT,VRA,VRB & \((\mathrm{Rc}=0)\) \\
vcmpgtsh. & VRT, VRA, VRB & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 04 & \[
\sigma_{6} \mathrm{VRT}
\] & \({ }_{11}\) VRA & \[
{ }_{16} \mathrm{VRB}
\] & RC \({ }_{21} 1_{22}\) & 838 & \\
\hline
\end{tabular}
```

do $\mathrm{i}=0$ to 127 by 16
$\operatorname{VRT}_{i: i+15} \leftarrow\left((\mathrm{VRA})_{i: i+15}>_{\text {si }}(\mathrm{VRB})_{i: i+15}\right) \quad{ }^{16} 1:{ }^{16} 0$
end
if $\mathrm{Rc}=1$ then do
$t \leftarrow\left(\right.$ VRT $\left.={ }^{128} 1\right)$
$\mathrm{f} \leftarrow\left(\mathrm{VRT}={ }^{128} 0\right)$
$\mathrm{CR6} \leftarrow \mathrm{t}\|\mathrm{ObO}\| \mathrm{f} \| \mathrm{ObO}$
end

```

For each integer value i from 0 to 7 , do the following. Signed-integer halfword element \(i\) in VRA is compared to signed-integer halfword element in VRB. Halfword element \(i\) in VRT is set to all 1s if signed-integer halfword element \(i\) in VRA is greater than signed-integer halfword element i in VRB, and is set to all Os otherwise.

\section*{Special Registers Altered:}

CR field 6
(if Rc=1)

\section*{Vector Compare Greater Than Signed Word VC-form}
\begin{tabular}{lll} 
vcmpgtsw & VRT,VRA,VRB & \((\mathrm{Rc}=0)\) \\
vcmpgtsw. & VRT \(, \mathrm{VRA}, \mathrm{VRB}\) & \((\mathrm{Rc}=1)\)
\end{tabular}

\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 32 \\
& \quad \operatorname{VRT}_{i: i+31} \leftarrow\left((\mathrm{VRA})_{i: i+31}>_{\text {si }}(\mathrm{VRB})_{i: i+31}\right) ?{ }^{32} 1:{ }^{32} 0 \\
& \text { end } \\
& \text { if Rc=1 then do } \\
& \quad \mathrm{t} \leftarrow\left(\mathrm{VRT}={ }^{128} 1\right) \\
& \quad \mathrm{f} \leftarrow\left(\mathrm{VRT}={ }^{128} 0\right) \\
& \quad \mathrm{CR6} \leftarrow \mathrm{t}\|0 \mathrm{bo}\| \mathrm{f} \| 0 \mathrm{bo} \\
& \text { end }
\end{aligned}
\]

For each integer value i from 0 to 3 , do the following. Signed-integer word element i in VRA is compared to signed-integer word element \(i\) in VRB. Word element \(i\) in VRT is set to all 1s if signed-integer word element \(i\) in VRA is greater than signed-integer word element i in VRB, and is set to all Os otherwise.

Special Registers Altered:
CR field 6
(if Rc=1)

\section*{Vector Compare Greater Than Unsigned Byte VC-form}
\begin{tabular}{lll} 
vcmpgtub & VRT,VRA,VRB & \((R c=0)\) \\
vcmpgtub. & VRT,VRA,VRB & \((R c=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & VRA & VRB & Rc & 518 & \\
\hline 0 & & 6 & & 11 & & & \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 8 \\
& \operatorname{VRT}_{i: i+7} \leftarrow\left((\mathrm{VRA})_{i: i+7}>_{\text {ui }}(\mathrm{VRB})_{i: i+7}\right) \quad{ }^{8}{ }^{8}:{ }^{8} 0 \\
& \text { end } \\
& \text { if } \mathrm{Rc}=1 \text { then do } \\
& \mathrm{t} \leftarrow\left(\mathrm{VRT}={ }^{128} 1\right) \\
& \mathrm{f} \leftarrow\left(\mathrm{VRT}={ }^{128} 0\right) \\
& \mathrm{CR6} \leftarrow \mathrm{t}\|\mathrm{ObO}|\mid \mathrm{f} \| \mathrm{ObO} \\
& \text { end }
\end{aligned}
\]

For each integer value i from 0 to 15 , do the following. Unsigned-integer byte element i in VRA is compared to unsigned-integer byte element \(i\) in VRB. Byte element i in VRT is set to all 1s if unsigned-integer byte element \(i\) in VRA is greater than to unsigned-integer byte element \(i\) in VRB, and is set to all 0s otherwise.

\section*{Special Registers Altered:}

CR field 6 (if \(\mathrm{Rc}=1\) )

\section*{Vector Compare Greater Than Unsigned Doubleword VX-form}
\begin{tabular}{lll} 
vcmpgtud & VRT,VRA, VRB & \((R c=0)\) \\
vcmpgtud. & VRT,VRA,VRB & \((R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & RC & 711 & \\
\hline 0 & & 61 & & & & \\
\hline
\end{tabular}
```

do i = 0 to 1
aop \leftarrow EXTZ (VR[VRA] .dword[i])
bop \& EXTZ(VR[VRB].dword[i])
if (EXTZ(aop) >ui EXTZ(bop)) then do
VR[VRT].dword[i] \leftarrow 0xFFFF_FFFF_FFFF_FFFF
flag.bit[i] \leftarrow0b1
end
else do
VR[VRT].dword[i] \leftarrow0x0000_0000_0000_0000
flag.bit[i] \leftarrow 0b1
end
end
if "vcmpgtud." then do
CR.bit[24]}\leftarrow(flag=0b11
CR.bit[25] \leftarrow 0b0
CR.bit[26] \leftarrow(flag=0b00)
CR.bit[27] \leftarrow0b0
end

```

For each integer value i from 0 to 1, do the following.
The unsigned integer value in doubleword element \(i\) of VR[VRA] is compared to the unsigned integer value in doubleword element \(i\) of VR[VRB]. Doubleword element \(i\) of VR[VRT] is set to all 1 s if the unsigned integer value in doubleword element \(i\) of VR[VRA] is greater than the unsigned integer value in doubleword element \(i\) of VR[VRB], and is set to all 0s otherwise.

Special Registers Altered:
CR field 6
(if \(R c=1\) )

\section*{Vector Compare Greater Than Unsigned Halfword VC-form}
\begin{tabular}{lll} 
vcmpgtuh & VRT,VRA,VRB & \((\mathrm{Rc}=0)\) \\
vcmpgtuh. & VRT,VRA,VRB & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & Rc & & 582 & \\
\hline 0 & & 61 & & 11 & & & \\
\hline
\end{tabular}
```

do $\mathrm{i}=0$ to 127 by 16
$\operatorname{VRT}_{i: i+15} \leftarrow\left((\operatorname{VRA})_{i: i+15}>_{\text {ui }}(\mathrm{VRB})_{i: i+15}\right) \quad{ }^{16} 1:{ }^{16} 0$
end
if $\mathrm{Rc}=1$ then do
$t \leftarrow\left(\right.$ VRT $\left.={ }^{128} 1\right)$
$\mathrm{f} \leftarrow\left(\mathrm{VRT}={ }^{128} 0\right)$
$\mathrm{CR} 6 \leftarrow \mathrm{t}||\mathrm{ObO}|| \mathrm{f} \| \mathrm{||}$ 0b0
end

```

For each integer value i from 0 to 7, do the following. Unsigned-integer halfword element i in VRA is compared to unsigned-integer halfword element i in VRB. Halfword element i in VRT is set to all 1s if unsigned-integer halfword element \(i\) in VRA is greater than to unsigned-integer halfword element i in VRB, and is set to all Os otherwise.

\section*{Special Registers Altered:}

CR field 6
(if Rc=1)

\section*{Vector Compare Greater Than Unsigned Word VC-form}
\begin{tabular}{lll} 
vcmpgtuw & VRT,VRA,VRB & \((\mathrm{Rc}=0)\) \\
vcmpgtuw. & VRT \(, \mathrm{VRA}, \mathrm{VRB}\) & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & RC & 646 & \\
\hline 0 & & 6 & & & 11 & & & \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 32 \\
& \quad \operatorname{VRT}_{i}: i+31 \leftarrow\left((\mathrm{VRA})_{i: i+31}>_{\text {ui }}(\mathrm{VRB})_{i: i+31}\right) \quad{ }^{32} 1:{ }^{32} 0 \\
& \text { end } \\
& \text { if Rc=1 then do } \\
& \quad \mathrm{t} \leftarrow\left(\mathrm{VRT}={ }^{128} 1\right) \\
& \quad \mathrm{f} \leftarrow\left(\mathrm{VRT}={ }^{128} 0\right) \\
& \quad \mathrm{CR6} 6 \mathrm{t}\|0 \mathrm{O} 0\| \mathrm{f} \| 0 \mathrm{obO} \\
& \text { end }
\end{aligned}
\]

For each integer value i from 0 to 3 , do the following. Unsigned-integer word element \(i\) in VRA is compared to unsigned-integer word element \(i\) in VRB. Word element \(i\) in VRT is set to all 1s if unsigned-integer word element i in VRA is greater than to unsigned-integer word element i in VRB, and is set to all 0s otherwise.

\section*{Special Registers Altered:}

CR field 6
(if Rc=1)

```

if MSR.VEC=0 then Vector_Unavailable()
for i = 0 to 15
srcl }\leftarrowVR[VRA],byte[i
src2 \leftarrow VR[VRB].byte[i]
if (srcl!= src2) then
VR[VRT],byte[i]}\leftarrow0xF
else
VR[VRT], byte[i] <0x00
end
al| true \leftarrow (VR[VRT]=OXFFFF FFFF FFFF FFFF FFF FFFF FFFF FFFF)

```

```

if RC=1 then CR, bit[56:59] \leftarrow(al|_true<<3) + (al| false<<l)

```

For each integer value i from 0 to 15 , do the following. The integer value in byte element \(i\) in VR[VRA] is compared to the integer value in byte element \(i\) in VR[ VRB] . The contents of byte element \(i\) in VR[ VRT] are set to \(0 \times F F\) if integer value in byte element \(i\) in VR[VRA] is not equal to the integer value in byte element \(i\) in VR[VRB], and are set to \(0 \times 00\) otherwise.

If \(R C=1, C R\) field 6 is set to indicate whether all vector elements compared true and whether all vector elements compared false.

\section*{Special Registers Altered:}

CR field 6
(if \(R C=1\) )

\section*{Vector Compare Not Equal or Zero Byte VX-form}
\begin{tabular}{lll} 
vcmpnezb & VRT,VRA,VRB & (if \(R c=0\) ) \\
vcmpnezb. & \(V R T, V R A, V R B\) & (if \(R c=1\) )
\end{tabular}
\begin{tabular}{|c|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & RC & 263 & \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

if MSR.VEC=O then Vector_Unavailable(l
for i = 0 to 15
srcl \leftarrowVR[VRA].byte[i]
srC2 \&VR[VRB],byte[i]
if (srcl = 0) | (src) = 0) | (srcl ! = src2) then
VR[VRT], byte[i] < OxFF
else
VR[VRT], byte[i] < 0x00
end
all_true \leftarrow(VR[VRT]=OxFFFF_FFFF_FFFF_FFF_FFF_FFF_FFFF_FFFF)
all_false}\leftarrow(VR[VRT)=0\times0000-0000-0000-0000-0000-0000_0000 - 0000)
if RC=1 then CR. bit[56:59] \leftarrow(al|_true<<3) + (a| _false<<1)

```

For each integer value i from 0 to 15 , do the following. The integer value in byte element \(i\) in VR[VRA] is compared to the integer value in byte element \(i\) in VR[ VRB]. The contents of byte element \(i\) in VR[ VRT] are set to \(0 \times F F\) if integer value in byte element \(i\) in VR[VRA] is not equal to the integer value in byte element i in VR[VRB] or either value is equal to \(0 \times 00\), and are set to \(0 \times 00\) otherwise.

If \(R C=1, C R\) field 6 is set to indicate whether all vector elements compared true and whether all vector elements compared false.

Special Registers Altered:
CR field 6
(if \(R c=1\) )

\section*{Vector Compare Not Equal Halfword VX-form}
\begin{tabular}{lll} 
vcmpneh & VRT, VRA,VRB & (if \(R c=0\) ) \\
vcmpneh. & VRT,VRA,VRB & (if \(R c=1\) )
\end{tabular}
\begin{tabular}{|c|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & Rc & 71 & \\
\hline 0 & & & & & & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
for i = 0 to 7
srcl \leftarrowVR[VRA].hword[i]
sIC2 \& VR[VRB], hword[i]
if (srcl!= src2) then
VR[ VRT].hword[i] < OxFFFF
else
VR[ VRT].hword[i] ↔0x0000
end
all_true \& (VR[VRT]=OxFFFF_FFFF_FFFF_fFFF_FF_ FFFF_FFFF_FFFF)

```

```

if Rc=1 then CR.bit[56:59] \& (al|_true<<3) + (a| I_false<<1)

```

For each integer value \(i\) from 0 to 7 , do the following. The integer value in halfword element \(i\) in VR[ VRA] is compared to the integer value in halfword element i in VR[VRB]. The contents of halfword element i in VR[VRT] are set to 0xFFFF if integer value in halfword element \(i\) in VR[VRA] is not equal to the integer value in halfword element i in VR[VRB] , and are set to \(0 \times 0000\) otherwise.

If \(R C=1, C R\) field 6 is set to indicate whether all vector elements compared true and whether all vector elements compared false.

\section*{Special Registers Altered:}

CR field \(6 \quad\) (if \(R C=1\) )

\section*{Vector Compare Not Equal or Zero Halfword VX-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{vcmpnezh vcmpnezh.}} & \multicolumn{2}{|l|}{\multirow[t]{2}{*}{VRT,VRA,VRB VRT,VRA,VRB}} & & \multicolumn{2}{|r|}{\multirow[t]{2}{*}{\[
\begin{aligned}
& \text { (if } R C=0 \text { ) } \\
& \text { (if } R C=1 \text { ) }
\end{aligned}
\]}} \\
\hline & & & & & & \\
\hline 4 & VRT & 11 VRA & \({ }_{16}\) VRB & \(\left\lvert\, \begin{aligned} & \text { RC } \\ & 21\end{aligned}\right.\) & 327 & 31 \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
for i = 0 to 7
srcl \leftarrowVR[VRA],hword[i]
src2 \leftarrowVR[VRB], hword[i]
if (srcl=0) | (src2=0) | (srcl \#src2) then
VR[VRT].hword[i] \& OxFFFF
else
VR[VRT].hword[i] <0x0000
end
al|_true \leftarrow(VR[VRT]=OXFFFFFFFF_FFFF_FFFF_FFFFFFFFFFFFFFFF)
all_false \leftarrow(VR[VRT] =0x0000-0000-0000-0000-0000_000-000-000000)
if RC=1 then CR, bit[56:59] \leftarrow(al|_true<<3) + (all_false<<l)

```

For each integer value i from 0 to 7, do the following. The integer value in halfword element \(i\) in VR[ VRA] is compared to the integer value in halfword element i in VR[VRB]. The contents of halfword element \(i\) in VR[VRT] are set to OXFFFF if integer value in halfword element \(i\) in VR[VRA] is not equal to the integer value in halfword element i in VR[ VRB] or either value is equal to \(0 \times 0000\), and are set to \(0 \times 0000\) otherwise.

If \(R C=1, C R\) field 6 is set to indicate whether all vector elements compared true and whether all vector elements compared false.

Special Registers Altered:
CR field 6
(if \(R C=1\) )

```

if MSR.VEC=0 then Vector_Unavailable()
for i = 0 to 3
srcl }\leftarrowVR[VRA].word[i
src2 \leftarrow VR[VRB].word[i]
if (srcl!= src2) then
VR[VRT],word[i] \&OxFFFF_FFFF
else
VR[VRT].word[i] }\leftarrow0\times0000_000
end
al| true \leftarrow(VR[VRT]=OXFFFF FFFF FFFF FFFF_FFF FFFF FFFF FFFF)

```

```

if Rc=1 then CR, bit[56:59] \leftarrow(al|_true<<3) + (a||_false<<l)

```

For each integer value i from 0 to 3 , do the following.
The integer value in word element \(i\) in VR[VRA] is compared to the integer value in word element \(i\) in VR[VRB]. The contents of word element i in VR[VRT] are set to OXFFFF_FFFF if integer value in word element i in VR[VRA] is not equal to the integer value in word element \(i\) in VR[VRB], and are set to \(0 \times 0000 \_0000\) otherwise.

If \(R C=1, C R\) field 6 is set to indicate whether all vector elements compared true and whether all vector elements compared false.

\section*{Special Registers Altered:}

CR field 6
(if \(R C=1\) )

\section*{Vector Compare Not Equal or Zero Word VX-form}
\begin{tabular}{lll} 
vcmpnezw & VRT,VRA,VRB & (if \(R c=0\) ) \\
vcmpnezw. & VRT,VRA,VRB & (if \(R c=1\) )
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & RC & 391 \\
\hline 0 & & & & & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
for i = 0 to 3
srcl \leftarrowVR[VRA].word[i]
srcl \leftarrowVR[VRB].word[i]
if (srcl=0) | (src2=0) | (srcl\not=src2) then
VR[VRT],word[i] < OXFFFF_FFFF
else
VR[VRT],word[i] }\leftarrow0\times0000_000
end
all_true \leftarrow(VR[VRT]=OXFFFF_FFFFFFFF_FFFFFFF_FFFF_FFFF_FFFF)
all_false \leftarrow (VR[VRT] =0\times0000-0000-0000_0000-0000 0000 0000-0000)
if RC=1 then CR, bit[56:59] \leftarrow(al|_true<<3) + (al|_false<<1)

```

For each integer value i from 0 to 3 , do the following. The integer value in word element \(i\) in VR[VRA] is compared to the integer value in word element \(i\) in VR[VRB]. The contents of word element i in VR[VRT] are set to OXFFFF_FFFF if integer value in word element \(i\) in VR[VRA] is not equal to the integer value in word element \(i\) in VR[VRB] or either value is equal to \(0 \times 0000 \_0000\), and are set to \(0 \times 0000 \_0000\) otherwise.

If \(R C=1, C R\) field 6 is set to indicate whether all vector elements compared true and whether all vector elements compared false.

\section*{Special Registers Altered:}

CR field 6
(if \(R c=1\) )

\subsection*{6.9.4 Vector Logical Instructions}

\section*{Extended mnemonics for vector logical operations}

Extended mnemonics are provided that use the Vector OR and Vector NOR instructions to copy the contents of one Vector Register to another, with and without complementing. These are shown as examples with the two instructions.

Vector Move Register
Several vector instructions can be coded in a way such that they simply copy the contents of one Vector Register to another. An extended mnemonic is provided to convey the idea that no computation is being performed but merely data movement (from one register to another).

The following instruction copies the contents of register Vy to register Vx.
vmr \(\mathrm{Vx}, \mathrm{Vy}\) (equivalent to: vor \(\mathrm{Vx}, \mathrm{Vy}, \mathrm{Vy}\) )

\section*{Vector Complement Register}

The Vector NOR instruction can be coded in a way such that it complements the contents of one Vector Register and places the result into another Vector Register. An extended mnemonic is provided that allows this operation to be coded easily.

The following instruction complements the contents of register Vy and places the result into register Vx.
vnot \(\mathrm{Vx}, \mathrm{Vy}\) (equivalent to: vnor \(\mathrm{Vx}, \mathrm{Vy}, \mathrm{Vy}\) )

\section*{Vector Logical AND VX-form}
Vand
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & VRT, VRA, VRB \\
\hline 0 & & 6 & & VRA & VRB \\
\hline 11 & & 1028 & \\
\hline
\end{tabular}
```

VR[VRT] \leftarrowVR[VRA] \& VR[VRB]

```

The contents of VR[VRA] are ANDed with the contents of VR[ VRB] and the result is placed into VR[ VRT].

\section*{Special Registers Altered: \\ None}

Vector Logical AND with Complement VX-form
vandc VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|ccc|}
\hline 4 & VRT & VRA & VRB & & 1092 & \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
\[
\text { VR[ VRT] } \leftarrow \text { VR[ VRA] \& } \operatorname{VR}[\text { VRB] }
\]

The contents of VR[VRA] are ANDed with the complement of the contents of VR[VRB] and the result is placed into VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Logical Equivalence VX-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{veqv VRT,VRA,VRB} \\
\hline 04 & \({ }_{6}\) VRT & VRA & VRB & 21 & 1668 \\
\hline
\end{tabular}
\[
\text { VR[ VRT] } \leftarrow V R[V R A] \equiv V R[V R B]
\]

The contents of VR[VRA] are XORed with the contents of VR[VRB] and the complemented result is placed into VR[ VRT].

\section*{Special Registers Altered: \\ None}

\section*{Vector Logical NAND VX-form}
vnand VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|ccc|}
\hline 4 & VRT & VRA & VRB & & 1412 & \\
\hline 0 & & 61 \\
\hline
\end{tabular}
```

if MSR,VEC=0 then VECTOR_UNAVAILABLE()
VR[VRT] \leftarrow\neg( VR[VRA] \& VR[VRB] )

```

The contents of VR[VRA] are ANDed with the contents of VR[ VRB] and the complemented result is placed into VR[ VRT].
```

Special Registers Altered:
None

```

\section*{Vector Logical OR with Complement VX-form}
vorc VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1348} \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
\[
\text { VR[VRT] } \leftarrow \text { VR[VRA] | -VR[VRB] }
\]

The contents of VR[VRA] are ORed with the complement of the contents of VR[VRB] and the result is placed into VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Logical NOR VX-form}
vnor VRT,VRA,VRB


The contents of VR[VRA] are ORed with the contents of VR[VRB] and the complemented result is placed into VR[VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Logical OR VX-form}
vor VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|ccc|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{1156} & \\
\hline 0 & & 61 \\
\hline
\end{tabular}
\[
\text { VR[VRT] } \leftarrow V R[V R A] \text { | VR[VRB] }
\]

The contents of VR[VRA] are ORed with the contents of VR[VRB] and the result is placed into VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Logical XOR VX-form}
vxor VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & & 1220 & \\
\hline 0 & & 6 & & 11 \\
\hline
\end{tabular}
\[
V R[V R T] \leftarrow V R[V R A] \oplus V R[V R B]
\]

The contents of VR[VRA] are XORed with the contents of VR[ VRB] and the result is placed into VR[ VRT].

\section*{Special Registers Altered:}

None

\subsection*{6.9.5 Vector Parity Byte Instructions}
```

Vector Parity Byte Word VX-form
vprtybw VRT,VRB

| 4 | VRT | 8 |  | VRB |  | 1538 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |  |

```
```

if MSR.VEC=O then Vector Unavailable()

```
if MSR.VEC=O then Vector Unavailable()
do \(\mathrm{i}=0\) to 3
do \(\mathrm{i}=0\) to 3
    \(s \leftarrow 0\)
    \(s \leftarrow 0\)
    do \(j=0\) to 3
    do \(j=0\) to 3
        \(s \leftarrow s\) ^ VR[VRB], word[i], byte[j], bit[7]
        \(s \leftarrow s\) ^ VR[VRB], word[i], byte[j], bit[7]
    end
    end
    VR[ VRT]. word[i] \(\leftarrow\) Chopl EXTZ(s), 32|
    VR[ VRT]. word[i] \(\leftarrow\) Chopl EXTZ(s), 32|
end
```

end

```

For each integer value \(i\) from 0 to 3 , do the following If the sum of the least significant bit in each byte sub-element of word element \(i\) of VR[VRB] is odd, the value 1 is placed into word element i of VR[ VRT] ; otherwise the value 0 is placed into word element \(i\) of VR[ VRT].

Special Registers Altered:
None
```

Vector Parity Byte Doubleword VX-form
vprtybd VRT,VRB

| 4 | VRT | 9 |  | VRB |  | 1538 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  |  |

if MSR.VEC=O then Vector_Unavailablell
do $i=0$ to 1
$s \leftarrow 0$
do $j=0$ to 7
$s \leftarrow s^{\wedge}$ VR[VRB]dword[i]. byte[j]. bit[7]
end
VR[VRT],dword[i] $\leftarrow$ Chop(EXTZ(s), 64)
end

```

For each integer value i from 0 to 1, do the following If the sum of the least significant bit in each byte sub-element of doubleword element i of VR[VRB] is odd, the value 1 is placed into doubleword element i of VR[VRT]; otherwise the value 0 is placed into doubleword element \(i\) of VR[ VRT].

Special Registers Altered:
None

\section*{Vector Parity Byte Quadword VX-form}
vprtybq VRT,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & 6 & VRT & 110 & VRB & \\
\hline 0 & 1538 & & \\
\hline
\end{tabular}
```

if MSR. VEC=O then Vector_Unavailable()
$s \leftarrow 0$
do $j=0$ to 15
$s \leftarrow s$ ^VR[VRB], byte[j], bit[7]
end
VR[VRT] $\leftarrow$ Chopl $\operatorname{EXTZ}(s), 128)$

```

If the sum of the least significant bit in each byte element of VR[VRB] is odd, the value 1 is placed into VR[ VRT] ; otherwise the value 0 is placed into VR[ VRT].

Special Registers Altered:
None

\subsection*{6.9.6 Vector Integer Rotate and Shift Instructions}

\section*{Vector Rotate Left Byte VX-form}
Vrlb
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & VRT,VRA,VRB \\
\hline 0 & & 6 & VRA & VRB & \\
\hline 11 & & 4 & 4 & 31 \\
\hline
\end{tabular}
```

do i=0 to 127 by }
sh}\leftarrow(VRB\mp@subsup{)}{i+5:i+7}{
VRT
end

```

For each integer value i from 0 to 15 , do the following. Byte element i in VRA is rotated left by the number of bits specified in the low-order 3 bits of the corresponding byte element \(i\) in VRB.

The result is placed into byte element i in VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Rotate Left Halfword VX-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|l|}{vrlh VRT,VRA,VRB} \\
\hline 04 & \[
\sigma_{6} \text { VRT }
\] & 11 VRA & \({ }_{16}\) VRB & 21 & 68 & 31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 16 \\
& \qquad \operatorname{sh} \leftarrow(V R B)_{i+12: i+15} \\
& \qquad V R T_{i: i+15} \leftarrow(V R A)_{i: i+15} \lll \text { sh } \\
& \text { end }
\end{aligned}
\]

For each integer value ifrom 0 to 7 , do the following. Halfword element i in VRA is rotated left by the number of bits specified in the low-order 4 bits of the corresponding halfword element in VRB.

The result is placed into halfword element i in VRT.

Special Registers Altered:
None

\section*{Vector Rotate Left Word VX-form}
Vrlw
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & VRT,VRA, VRB \\
\hline 0 & & 6 & VRT & VRA & VRB \\
11 & & 132 & \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 32 \\
& \qquad \operatorname{sh} \leftarrow(\text { VRB })_{i+27: i+31} \\
& \qquad \operatorname{VRT}_{i: i+31} \leftarrow(\text { VRA })_{i: i+31} \lll \text { sh } \\
& \text { end }
\end{aligned}
\]

For each integer value ifrom 0 to 3 , do the following. Word element i in VRA is rotated left by the number of bits specified in the low-order 5 bits of the corresponding word element \(i\) in VRB.

The result is placed into word element i in VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Rotate Left Doubleword VX-form}
Vrld
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & VRT, VRA, VRB \\
\hline 0 & & 6 & VRA & VRB & \\
\hline 11 & & 196 & & & \\
\hline
\end{tabular}
```

do i = 0 to 1
sh}\leftarrowVR[VRB].dword[i].bit[58:63]
VR[VRT].dword[i] \leftarrow VR[VRA].dword[i] <<< sh
end

```

For each integer value i from 0 to 1 , do the following. The contents of doubleword element \(i\) of VR[ VRA] are rotated left by the number of bits specified in bits 58:63 of doubleword element \(i\) of VR[ VRB].

The result is placed into doubleword element i of VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Shift Left Byte VX-form}
\[
\text { vslb } \quad \text { VRT,VRA,VRB }
\]
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{260} \\
\hline 0 & & 61 \\
\hline
\end{tabular}
```

do i=0 to 127 by 8
sh}\leftarrow(VRB\mp@subsup{)}{i+5:i+7}{i
VRT i:i+7}\mp@code{\leftarrow(VRA) i:i+7}<<< s

```
end

For each integer value i from 0 to 15 , do the following. Byte element i in VRA is shifted left by the number of bits specified in the low-order 3 bits of byte element i in VRB.
- Bits shifted out of bit 0 are lost.
- Zeros are supplied to the vacated bits on the right.

The result is placed into byte element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Shift Left Halfword VX-form}
vslh VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 324 \\
\hline 0 & & 61 \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
sh}\leftarrow(\textrm{VRB}\mp@subsup{)}{i+12:i+15}{
VRT
end

```

For each integer value i from 0 to 7 , do the following. Halfword element i in VRA is shifted left by the number of bits specified in the low-order 4 bits of halfword element i in VRB.
- Bits shifted out of bit 0 are lost.
- Zeros are supplied to the vacated bits on the right.

The result is placed into halfword element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Shift Left Word VX-form}
vslw VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{388} \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 32 \\
& \qquad \operatorname{sh} \leftarrow(\mathrm{VRB})_{i+27: i+31} \\
& \qquad \operatorname{VRT}_{i: i+31} \leftarrow(\mathrm{VRA})_{i: i+31} \ll \text { sh } \\
& \text { end }
\end{aligned}
\]

For each integer value i from 0 to 3 , do the following. Word element \(i\) in VRA is shifted left by the number of bits specified in the low-order 5 bits of word element i in VRB.
- Bits shifted out of bit 0 are lost.
- Zeros are supplied to the vacated bits on the right.

The result is placed into word element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Shift Left Doubleword VX-form}
vsld VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{1476} \\
\hline 0 & & 611 & & & \\
\hline
\end{tabular}
```

do i = 0 to 1
sh}\leftarrowVR[VRB].dword[i].bit[58:63]
VR[VRT].dword[i] \leftarrow VR[VRA].dword[i] << sh
end

```

For each integer value i from 0 to 1, do the following.
The contents of doubleword element i of VR[ VRA]
are shifted left by the number of bits specified in bits 58:63 of doubleword element \(i\) of VR[ VRB].
- Bits shifted out of bit 0 are lost.
- Zeros are supplied to the vacated bits on the right.

The result is placed into doubleword element i of VR[VRT].

\section*{Special Registers Altered: \\ None}

\section*{Vector Shift Right Byte VX-form}
VRTb
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT, VRA, VRB \\
\hline 0 & & 6 & VRA & VRB & \\
\hline 11 & 516 & & \\
\hline
\end{tabular}

> do \(i=0\) to 127 by 8 \[ \operatorname{sh} \leftarrow(\mathrm{VRB})_{i+5: i+7} \] \(\quad \mathrm{VRT}_{i: i+7} \leftarrow(\mathrm{VRA})_{i: i+7} \gg_{\text {ui }}\) sh end

For each integer value i from 0 to 15 , do the following. Byte element i in VRA is shifted right by the number of bits specified in the low-order 3 bits of byte element \(i\) in VRB. Bits shifted out of the least-significant bit are lost. Zeros are supplied to the vacated bits on the left. The result is placed into byte element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Shift Right Halfword VX-form}
vsrh VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 580 \\
\hline 0 & & 61 & & & \\
\hline
\end{tabular}

> do \(i=0\) to 127 by 16 \[ \operatorname{sh} \leftarrow(\mathrm{VRB})_{i+12: i+15} \] \(\quad \operatorname{VRT}_{i: i+15} \leftarrow(\mathrm{VRA})_{i: i+15} \gg_{\text {ui }}\) sh end

For each integer value i from 0 to 7 , do the following. Halfword element i in VRA is shifted right by the number of bits specified in the low-order 4 bits of halfword element i in VRB. Bits shifted out of the least-significant bit are lost. Zeros are supplied to the vacated bits on the left. The result is placed into halfword element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Shift Right Word VX-form}
VRT,VRA, VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 644 \\
\hline 0 & & 6 & 11 & 16 & 21 \\
\hline
\end{tabular}

> do \(i=0\) to 127 by 32 \[ \operatorname{sh} \leftarrow(\mathrm{VRB})_{i+27: i+31} \] \[ \text { VRT }_{i: i+31} \leftarrow(\mathrm{VRA})_{i: i+31} \gg_{\text {ui }} \text { sh } \] end

For each integer value i from 0 to 3 , do the following. Word element \(i\) in VRA is shifted right by the number of bits specified in the low-order 5 bits of word element i in VRB. Bits shifted out of the least-significant bit are lost. Zeros are supplied to the vacated bits on the left. The result is placed into word element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Shift Right Doubleword VX-form}
vsrd VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & & 1732 \\
\hline 0 & & 61 & & & \\
\hline
\end{tabular}
```

do i = 0 to 1
sh}\leftarrowVR[VRB].dword[i].bit[58:63]
VR[VRT].dword[i] \leftarrow VR[VRA].dword[i] >> ui sh
end

```

For each integer value i from 0 to 1, do the following. The contents of doubleword element \(i\) of VR[ VRA] are shifted right by the number of bits specified in bits 58:63 of doubleword element i of VR[VRB]. Zeros are supplied to the vacated bits on the left.

The result is placed into doubleword element i of VR[ VRT].

Special Registers Altered:
None

\section*{Vector Shift Right Algebraic Byte VX-form}


For each integer value i from 0 to 15 , do the following. Byte element i in VRA is shifted right by the number of bits specified in the low-order 3 bits of the corresponding byte element i in VRB. Bits shifted out of bit 7 of the byte element are lost. Bit 0 of the byte element is replicated to fill the vacated bits on the left. The result is placed into byte element i of VRT .

\section*{Special Registers Altered:}

None

\section*{Vector Shift Right Algebraic Halfword VX-form}
Vsrah
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & \(0_{6}\) VRT & & VRA & VRB & \\
\hline 0 & & & 16 & 836 & \\
\hline
\end{tabular}
```

do i=0 to 127 by 16
sh}\leftarrow(\textrm{VRB}\mp@subsup{)}{i+12:i+15}{
VRT}\mp@subsup{T}{i:i+15}{*}\leftarrow(\textrm{VRA}\mp@subsup{)}{i:i+15}{}>>>>>\mp@code{si
end

```

For each integer value i from 0 to 7 , do the following. Halfword element i in VRA is shifted right by the number of bits specified in the low-order 4 bits of the corresponding halfword element i in VRB. Bits shifted out of bit 15 of the halfword are lost. Bit 0 of the halfword is replicated to fill the vacated bits on the left. The result is placed into halfword element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Shift Right Algebraic Word VX-form}
```

vsraw VRT,VRA,VRB

```
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{900} \\
\hline 0 & & 6 & & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
sh}\leftarrow(VRB\mp@subsup{)}{i+27:i+31}{

```

```

end

```

For each integer value ifrom 0 to 3 , do the following. Word element \(i\) in VRA is shifted right by the number of bits specified in the low-order 5 bits of the corresponding word element i in VRB. Bits shifted out of bit 31 of the word are lost. Bit 0 of the word is replicated to fill the vacated bits on the left. The result is placed into word element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Shift Right Algebraic Doubleword VX-form}
vsrad VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{964} \\
\hline 0 & & 6 & & & & \\
\hline
\end{tabular}
```

do i = 0 to 1
sh \leftarrow VR[VRB].dword[i].bit[58:63]
VR[VRT].dword[i] \leftarrow VR[VRA].dword[i] >> si sh
end

```

For each integer value ifrom 0 to 1 , do the following.
The contents of doubleword element i of VR[ VRA] are shifted right by the number of bits specified in bits \(58: 63\) of doubleword element \(i\) of VR[ VRB]. Bit 0 of doubleword element \(i\) of VR[VRA] is replicated to fill the vacated bits on the left.

The result is placed into doubleword element i of VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Rotate Left Word then AND with Mask VX-form}
Vrlwnm
\begin{tabular}{|l|l|l|l|l|lll|}
\hline 4 & VRT,VRA,VRB \\
\hline 0 & & 6 & VRT & VRA & VRB & & 389 \\
\hline 11 & & 16 & 21 & & 31 \\
\hline
\end{tabular}
if MSR.VEC=0 then Vector_Unavailablel)
```

do i = 0 to 3
srcl.word[0]}\leftarrowVR[VRA],word[i
srcl.word[1] \& VR[VRA],word[i]
srCl }\leftarrowVR[VRB],word[i
b
e}\leftarrow\operatorname{src2,bit[19:23]
n}\leftarrow\operatorname{src2.bit[27:31]
r}\leftarrow\operatorname{srcl,bit[n:n+31]
m}\leftarrow\operatorname{MASK}(b,e
VR[VRT], word[i] \leftarrowr\& m
end

```

For each integer value i from 0 to 3 , do the following.
Let srcl be the contents of word element i of VR[VRA].

Let \(\operatorname{src} 2\) be the contents of word element i of VR[ VRB].

Let mb be the contents of bits \(11: 15\) of src .
Let me be the contents of bits 19:23 of src .
Let \(s h\) be the contents of bits 27:31 of src2.
srcl is rotated left \(s h\) bits. A mask is generated having 1-bits from bit mb through bit me and 0-bits elsewhere. The rotated data are ANDed with the generated mask.

The result is placed into word element i of VR[VRT].

\section*{Special Registers Altered: None}

\section*{Vector Rotate Left Word then Mask Insert VX-form}
vrlwmi VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & \multicolumn{1}{|c|}{ VRB } & & 133 \\
\hline 0 & & 6 & & & & \\
\hline
\end{tabular}
if MSR.VEC=O then Vector _Unavailablell
```

do i = 0 to 3
srcl.word[0]}\leftarrow\mathrm{ VR[VRA].word[i]
srcl.word[1] \leftarrowVR[VRA], word[i]
srcl }\leftarrowVR[VRB], word[i
src3 \&VR[VRT], word[i]
b}\leftarrow\mathrm{ src2.bit[11:15]
e}\leftarrow\mathrm{ src2.bit[19:23]
n}\leftarrow\mathrm{ src2,bit[27:31]
r}\leftarrow\mathrm{ srcl,bit[n:n+31]
m}\leftarrow\operatorname{MASK}(b,e
VR[VRT],word[i] \& (r\&m)| (src3\& \&m)
end

```

For each integer value i from 0 to 3 , do the following. Let sicl be the contents of word element i of VR[ VRA].

Let \(\operatorname{src} 2\) be the contents of word element i of VR[ VRB].

Let \(\operatorname{srcs}\) be the contents of word element i of VR[VRT].

Let mb be the contents of bits \(11: 15\) of src .
Let me be the contents of bits \(19: 23\) of src 2 .
Let \(\mathrm{s} h\) be the contents of bits \(27: 31\) of src .
srcl is rotated left sh bits. A mask is generated having 1-bits from bit mb through bit me and 0-bits elsewhere. The rotated data are inserted into srcs under control of the generated mask.

The result is placed into word element i of VR[ VRT].

\section*{Special Registers Altered: \\ None}

\section*{Vector Rotate Left Doubleword then AND with Mask VX-form}
vrldnm VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & & 453 & 31 \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailable()
```

do i = 0 to 1
srcl.dword[0] \leftarrow VR[VRA].dword[i]
srcl.dword[1]}\leftarrowVR[VRA].dword[i]
srcl \leftarrowVR[VRB].dword[i]
b}\leftarrow\operatorname{src2.bit[42:47]
e}\leftarrow\mathrm{ src2.bit[50:55]
n}\leftarrow\operatorname{src2.bit[58:63]
r}\leftarrow\mathrm{ srcl.bit[n:n+63]
m}\leftarrow\operatorname{MASK}(b,e
VR[VRT].dword[i] <r\&m
end

```

For each integer value i from 0 to 1, do the following. Let srcl be the contents of doubleword element of VR[ VRA].

Let src2 be the contents of doubleword element i of VR[ VRB].

Let mb be the contents of bits \(42: 47\) of \(\mathrm{src2}\).
Let me be the contents of bits 50 : 55 of \(\mathrm{src2}\). Let \(s h\) be the contents of bits \(58: 63\) of \(\operatorname{src} 2\).
srcl is rotated left sh bits. A mask is generated having 1-bits from bit mb through bit me and 0-bits elsewhere. The rotated data are ANDed with the generated mask.

The result is placed into doubleword element \(i\) of VR[ VRT].

Special Registers Altered:
None

\section*{Vector Rotate Left Doubleword then Mask Insert VX-form}
vrldmi VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|ccc|}
\hline 4 & 6 & VRT & VRA & VRB & & 197 \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailable()
do \(i=0\) to 1
sicl.dword[0] \(\leftarrow\) VR[VRa].dword[i]
sccl.dword[1] \(\leftarrow\) VR[VRa]. dword[i]
sicl \(\leftarrow\) VR[ VRB]. dword[i]
\(\operatorname{src} 3 \leftarrow\) VR[ VRT]. dword \([i]\)
\(b \leftarrow\) sich. bit [42:47]
\(e \leftarrow \operatorname{sic} 2\), bit [ \([50: 55]\)
\(n \leftarrow\) sicl. bit [ 58 : 63]
\(1 \leftarrow\) sicl.bit [n:n+63]
\(m \leftarrow \operatorname{MASK}(b, e)\)
VR[ VRT]. dwordili \(\leftarrow(r \& m) \mid(\operatorname{sic} 3 \&-m)\)
end
For each integer value i from 0 to 1 , do the following. Let srcl be the contents of doubleword element \(i\) of VR[VRA].

Let src2 be the contents of doubleword element i of VR[ VRB] .

Let \(\operatorname{srcs}\) be the contents of doubleword element i of VR[ VRT] .

Let mb be the contents of bits \(42: 47\) of \(\mathrm{src2}\).
Let me be the contents of bits \(50: 55\) of \(\mathrm{src2}\).
Let sh be the contents of bits \(58: 63\) of src .
srcl is rotated left sh bits. A mask is generated having 1-bits from bit mb through bit me and 0-bits elsewhere. The rotated data are inserted into stc3 under control of the generated mask.

The result is placed into doubleword element i of VR[VRT].

Special Registers Altered:
None

\subsection*{6.10 Vector Floating-Point Instruction Set}

\subsection*{6.10.1 Vector Floating-Point Arithmetic Instructions}

\section*{Vector Add Floating-Point VX-form}
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & & 10 \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
VRT i:i+31
end

```

For each integer value ifrom 0 to 3 , do the following. Single-precision floating-point element \(i\) in VRA is added to single-precision floating-point element i in VRB. The intermediate result is rounded to the nearest single-precision floating-point number and placed into word element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Subtract Floating-Point VX-form}
vsubfp VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|ccc|}
\hline 4 & VRT & VRA & VRB & & 74 & \\
\hline 0 & & 61 \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
VRT i:i+31
end

```

For each integer value i from 0 to 3 , do the following. Single-precision floating-point element \(i\) in VRB is subtracted from single-precision floating-point element \(i\) in VRA. The intermediate result is rounded to the nearest single-precision floating-point number and placed into word element i of VRT.

Special Registers Altered:
None

\section*{Vector Multiply-Add Floating-Point VA-form}
```

vmaddfp VRT,VRA,VRC,VRB

```
\begin{tabular}{|l|l|l|l|l|ll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & VRC & \multicolumn{2}{|c|}{46} & \\
\hline 0 & & 6 & & 11 & & 26 \\
\hline
\end{tabular}
```

do $i=0$ to 127 by 32
$\operatorname{prod} \leftarrow(V R A)_{i: i+31} x_{f p}(V R C)_{i: i+31}$
$\operatorname{VRT}_{i: i+31} \leftarrow$ RoundToNearSP( prod $\left.+_{f p}(V R B)_{i: i+31}\right)$
end

```

For each integer value i from 0 to 3 , do the following. Single-precision floating-point element in VRA is multiplied by single-precision floating-point element i in VRC. Single-precision floating-point element i in VRB is added to the infinitely-precise product. The intermediate result is rounded to the nearest single-precision floating-point number and placed into word element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Programming Note}

To use a multiply-add to perform an IEEE or Java compliant multiply, the addend must be -0.0. This is necessary to insure that the sign of a zero result will be correct when the product is \(-0.0(+0.0+-0.0\) \(\geq+0.0\), and \(-0.0+-0.0 \geq-0.0)\). When the sign of a resulting 0.0 is not important, then +0.0 can be used as an addend which may, in some cases, avoid the need for a second register to hold a -0.0 in addition to the integer 0/floating-point +0.0 that may already be available.

\section*{Vector Negative Multiply-Subtract Floating-Point VA-form}
vnmsubfp VRT,VRA,VRC,VRB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & VRC & \multicolumn{2}{|c|}{47} \\
\hline 0 & & 6 & & & & \\
\hline
\end{tabular}
do \(\mathrm{i}=0\) to 127 by 32
\[
\begin{aligned}
& \operatorname{prod}_{0: i n f} \leftarrow(V R A)_{i: i+31} x_{f p}(V R C)_{i: i+31} \\
& \operatorname{VRT}_{i: i+31} \leftarrow-\text {-RoundToNearSP }\left(\operatorname{prod}_{0: i n f}-{ }_{f p}(V R B)_{i: i+31}\right)
\end{aligned}
\]
end

For each integer value i from 0 to 3 , do the following. Single-precision floating-point element \(i\) in VRA is multiplied by single-precision floating-point element i in VRC. Single-precision floating-point element \(i\) in VRB is subtracted from the infinitely-precise product. The intermediate result is rounded to the nearest single-precision floating-point number, then negated and placed into word element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\subsection*{6.10.2 Vector Floating-Point Maximum and Minimum Instructions}

\section*{Vector Maximum Floating-Point VX-form}
Vmaxfp
\begin{tabular}{|c|c|c|c|c|}
\hline 4 & VRT, VRA, VRB \\
\hline 0 & 6 & VRA & VRB & 1034 \\
\hline 16 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by }3
gt_flag}\leftarrow((VRA) i:i+31 > >fp (VRB) i:i+31 )
VRT
end

```

For each integer value i from 0 to 3 , do the following. Single-precision floating-point element \(i\) in VRA is compared to single-precision floating-point element i in VRB. The larger of the two values is placed into word element \(i\) of VRT.

The maximum of +0 and -0 is +0 . The maximum of any value and a NaN is a QNaN .

\section*{Special Registers Altered:}

None

Vector Minimum Floating-Point VX-form
VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|cc|}
\hline 4 & VRT & VRA & VRB & & 1098 \\
\hline 0 & & 6 & 11 & 16 & 21 \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
lt_flag}\leftarrow((VRA) i:i+31 < fp (VRB) i:i+31 )
VRT i:i+31
end

```

For each integer value i from 0 to 3 , do the following. Single-precision floating-point element \(i\) in VRA is compared to single-precision floating-point element i in VRB. The smaller of the two values is placed into word element \(i\) of VRT.

The minimum of +0 and -0 is -0 . The minimum of any value and a NaN is a QNaN .

Special Registers Altered:
None

\subsection*{6.10.3 Vector Floating-Point Rounding and Conversion Instructions}

See Appendix C, "Vector RTL Functions" on page 789, for RTL function descriptions.

\section*{Vector Convert To Signed Fixed-Point Word Saturate VX-form}

\author{
vctsxs VRT,VRB,UIM
}
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & UIM & VRB & & 970 \\
\hline 0 & & 61 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
VRT i:i+31}\leftarrow\leftarrow\mathrm{ ConvertSPtoSXWsaturate((VRB) i:i+31, UIM)
end

```

For each integer value i from 0 to 3 , do the following. Single-precision floating-point word element i in VRB is multiplied by \(2^{\mathrm{UIM}}\). The product is converted to a 32-bit signed fixed-point integer using the rounding mode Round toward Zero.
- If the intermediate result is greater than \(2^{31}-1\) the result saturates to \(2^{31}-1\).
- If the intermediate result is less than \(-2^{31}\) the result saturates to \(-2^{31}\).

The result is placed into word element \(i\) of VRT.
Special Registers Altered:

\section*{SAT}

\section*{Extended Mnemonics:}

Example of an extended mnemonics for Vector Convert to Signed Fixed-Point Word Saturate:
```

Extended: Equivalent to:
vcfpsxws VRT,VRB,UIM vctsxs VRT,VRB,UIM

```

\section*{Vector Convert To Unsigned Fixed-Point Word Saturate VX-form}
vctuxs VRT,VRB,UIM
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & UIM & VRB & & 906 \\
\hline 0 & & 61 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
VRT i:i+31
end

```

For each integer value i from 0 to 3 , do the following. Single-precision floating-point word element i in VRB is multiplied by \(2^{\mathrm{UIM}}\). The product is converted to a 32-bit unsigned fixed-point integer using the rounding mode Round toward Zero.
- If the intermediate result is greater than \(2^{32}-1\) the result saturates to \(2^{32}-1\).

The result is placed into word element \(i\) of VRT.

\section*{Special Registers Altered:} SAT

\section*{Extended Mnemonics:}

Example of an extended mnemonics for Vector Convert to Unsigned Fixed-Point Word Saturate:

\section*{Extended:}

Equivalent to:
vcfpuxws VRT,VRB,UIM vctuxs VRT,VRB,UIM

\section*{Vector Convert From Signed Fixed-Point Word VX-form}
vcfsx
VRT,VRB,UIM
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & UIM & VRB & & 842 & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
VRT}\mp@subsup{\textrm{i}:1+31}{}{~
end

```

For each integer value i from 0 to 3 , do the following. Signed fixed-point word element \(i\) in VRB is converted to the nearest single-precision floating-point value. Each result is divided by \(2^{\mathrm{UIM}}\) and placed into word element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Vector Convert from Signed Fixed-Point Word

\author{
Extended: \\ vcsxwfp VRT,VRB,UIM \\ Equivalent to: \\ vcfsx VRT,VRB,UIM
}

\section*{Vector Convert From Unsigned Fixed-Point Word VX-form}
vcfux VRT,VRB,UIM
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & UIM & VRB & & 778 & \\
\hline 0 & & 611 & & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
VRT i:i+31
end

```

For each integer value i from 0 to 3 , do the following. Unsigned fixed-point word element \(i\) in VRB is converted to the nearest single-precision floating-point value. The result is divided by \(2^{\text {UIM }}\) and placed into word element \(i\) of VRT.

\section*{Special Registers Altered:}

None

\section*{Extended Mnemonics:}

Examples of extended mnemonics for Vector Convert from Unsigned Fixed-Point Word

\author{
Extended: \\ Equivalent to: \\ vcuxwfp VRT,VRB,UIM vcfux VRT,VRB,UIM
}

\section*{Vector Round to Floating-Point Integer toward -Infinity VX-form}

\section*{vrfim}

VRT,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & & III & VRB & \\
\hline 0 & & 6 & & 714 & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
VRT 0:31}\leftarrow\leftarrow\mathrm{ RoundToSPIntFloor( (VRB) 0:31
end

```

For each integer value i from 0 to 3 , do the following. Single-precision floating-point element \(i\) in VRB is rounded to a single-precision floating-point integer using the rounding mode Round toward -Infinity.

The result is placed into the corresponding word element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Programming Note}

The Vector Convert To Fixed-Point Word instructions support only the rounding mode Round toward Zero. A floating-point number can be converted to a fixed-point integer using any of the other three rounding modes by executing the appropriate Vector Round to Floating-Point Integer instruction before the Vector Convert To Fixed-Point Word instruction.

\section*{Programming Note}

The fixed-point integers used by the Vector Convert instructions can be interpreted as consisting of 32-UIM integer bits followed by UIM fraction bits.

\section*{Vector Round to Floating-Point Integer Nearest VX-form}

\author{
vrfin VRT,VRB
}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & VRT & III & VRB & & 522 & \\
\hline 0 & 6 & 11 & 16 & 21 & & 31 \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
VRT}\mp@subsup{T}{0:31}{}\leftarrow\mathrm{ RoundToSPIntNear( (VRB) 0:31 )
end

```

For each integer value i from 0 to 3 , do the following. Single-precision floating-point element i in VRB is rounded to a single-precision floating-point integer using the rounding mode Round to Nearest.

The result is placed into the corresponding word element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Vector Round to Floating-Point Integer toward +Infinity VX-form}
vrfip
VRT,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & & III & VRB & \\
\hline 0 & & 650 & & & \\
\hline
\end{tabular}
do \(i=0\) to 127 by 32
\(\mathrm{VRT}_{0: 31} \leftarrow\) RoundToSPIntCeil( (VRB) \(0_{0: 31}\) )
end

For each integer value i from 0 to 3 , do the following. Single-precision floating-point element \(i\) in VRB is rounded to a single-precision floating-point integer using the rounding mode Round toward +Infinity.

The result is placed into the corresponding word element i of VRT.

\section*{Special Registers Altered:}

None

\section*{Version 3.0}

\section*{Vector Round to Floating-Point Integer toward Zero VX-form}
vrfiz
VRT,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & & III & VRB & \\
\hline 0 & & 686 & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
VRT}\mp@subsup{0}{0:31}{}\leftarrow\mathrm{ RoundToSPIntTrunc( (VRB) 0:31 )
end

```

For each integer value i from 0 to 3 , do the following. Single-precision floating-point element \(i\) in VRB is rounded to a single-precision floating-point integer using the rounding mode Round toward Zero.

The result is placed into the corresponding word element i of VRT.

Special Registers Altered:
None

\subsection*{6.10.4 Vector Floating-Point Compare Instructions}

The Vector Floating-Point Compare instructions compare two Vector Registers word element by word element, interpreting the elements as single-precision floating-point numbers. With the exception of the Vector Compare Bounds Floating-Point instruction, they set the target Vector Register, and CR Field 6 if Rc=1, in the same manner as do the Vector Integer Compare instructions; see Section 6.9.3.

\section*{Vector Compare Bounds Floating-Point VC-form}
\begin{tabular}{lll} 
vcmpbfp & VRT,VRA,VRB & \((\mathrm{Rc}=0)\) \\
vcmpbfp. & VRT,VRA,VRB & \((\mathrm{Rc}=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & \[
{ }_{6} \text { VRT }
\] & \[
11 \text { VRA }
\] & \begin{tabular}{l}
VRB \\
16
\end{tabular} & \(\left|\begin{array}{|c|}\mathrm{Rc} \\ 21\end{array}\right|_{22}\) & 966 & 31 \\
\hline
\end{tabular}
```

do $i=0$ to 127 by 32
$1 e \leftarrow\left((V R A)_{i: i+31} \leq_{f p}(V R B)_{i: i+31}\right)$
$g e \leftarrow\left((V R A)_{i: i+31} \geq_{\text {fp }}-(V R B)_{i: i+31}\right)$
$\forall R T_{i: i+31} \leftarrow$ ᄀle $\|$ ᄀge $\|^{30} 0$
end
if $\mathrm{Rc}=1$ then do
$\mathrm{ib} \leftarrow\left(\mathrm{VRT}={ }^{128} 0\right)$
$\mathrm{CR} 6 \leftarrow \mathrm{ObOO} \| \mathrm{ib}| | \mathrm{ObO}$
end

```

For each integer value i from 0 to 3 , do the following.
Single-precision floating-point word element i in VRA is compared to single-precision floating-point word element \(i\) in VRB. A 2-bit value is formed that indicates whether the element in VRA is within the bounds specified by the element in VRB, as follows.
- Bit 0 of the 2 -bit value is set to 0 if the element in VRA is less than or equal to the element in VRB, and is set to 1 otherwise.
- Bit 1 of the 2 -bit value is set to 0 if the element in VRA is greater than or equal to the negation of the element in VRB, and is set to 1 otherwise.

The 2-bit value is placed into the high-order two bits of word element \(i\) of VRT and the remaining bits of element \(i\) are set to 0 .

If \(R c=1, C R\) field 6 is set as follows.

\section*{Bit Description}

0 Set to 0
1 Set to 0

The Vector Compare Bounds Floating-Point instruction sets the target Vector Register, and CR Field 6 if Rc=1, to indicate whether the elements in VRA are within the bounds specified by the corresponding element in VRB, as explained in the instruction description. A single-precision floating-point value \(x\) is said to be "within the bounds" specified by a single-precision floating-point value \(y\) if \(-\mathrm{y} \leq \mathrm{x} \leq \mathrm{y}\).

\section*{Bit Description}

2 Set to indicate whether all four elements in VRA are within the bounds specified by the corresponding element in VRB, otherwise set to 0.

3 Set to 0

\section*{Special Registers Altered:}

CR field 6
(if \(\mathrm{Rc}=1\) )

\section*{Programming Note}

Each single-precision floating-point word element in VRB should be non-negative; if it is negative, the corresponding element in VRA will necessarily be out of bounds.

One exception to this is when the value of an element in VRB is -0.0 and the value of the corresponding element in VRA is either +0.0 or -0.0 . +0.0 and -0.0 compare equal to -0.0 .

\section*{Vector Compare Equal Floating-Point VC-form}
\begin{tabular}{lll} 
vcmpeqfp & VRT,VRA,VRB & \((R c=0)\) \\
vcmpeqfp. & VRT,VRA,VRB & \((R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 04 & \[
{ }_{6} \text { VRT }
\] & VRA & \[
{ }_{16} \text { VRB }
\] & RC 21. & 198 & 31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 32 \\
& \quad \operatorname{VRT}_{i: i+31} \leftarrow\left((\mathrm{VRA})_{i: i+31}={ }_{f p}(\mathrm{VRB})_{i: i+31}\right) ?{ }^{32} 1:{ }^{32} 0 \\
& \text { end } \\
& \text { if } R c=1 \text { then do } \\
& \quad \mathrm{t} \leftarrow(\mathrm{VRT}=1281) \\
& \quad \mathrm{f} \leftarrow\left(\mathrm{VRT}={ }^{128} 0\right) \\
& \quad \mathrm{CR6} 6 \mathrm{t}\|\mathrm{ObO}\| \mathrm{f} \| 0 \mathrm{~b} 0 \\
& \text { end }
\end{aligned}
\]

For each integer value i from 0 to 3 , do the following.
Single-precision floating-point element \(i\) in VRA is compared to single-precision floating-point element \(i\) in VRB. Word element \(i\) in VRT is set to all 1s if single-precision floating-point element \(i\) in VRA is equal to single-precision floating-point element i in VRB, and is set to all 0 s otherwise.

If the source element i in VRA or the source element \(i\) in VRB is a NaN, VRT is set to all Os, indicating "not equal to". If the source element \(i\) in VRA and the source element i in VRB are both infinity with the same sign, VRT is set to all 1 s , indicating "equal to".

\section*{Special Registers Altered:}

CR field 6 . (if Rc=1)

\section*{Vector Compare Greater Than or Equal Floating-Point VC-form}
\begin{tabular}{lll} 
vcmpgefp & VRT,VRA,VRB & \((R c=0)\) \\
vcmpgefp. & VRT,VRA,VRB & \((R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & Rc & 454 & \\
\hline 0 & 6 & 11 & 16 & 21.22 & & 31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 32 \\
& \quad V R T_{i: i+31} \leftarrow\left((V R A)_{i: i+31} \geq_{f p}(V R B)_{i: i+31}\right) ?^{32} 1:{ }^{32} 0 \\
& \text { end } \\
& \text { if } R c=1 \text { then do } \\
& \quad t \leftarrow\left(V R T={ }^{128} 1\right) \\
& \quad f \leftarrow\left(V R T={ }^{128} 0\right) \\
& \quad C R 6 \leftarrow t\|0 b 0\| f \| \text { Obo } \\
& \text { end }
\end{aligned}
\]

For each integer value i from 0 to 3 , do the following.
Single-precision floating-point element \(i\) in VRA is compared to single-precision floating-point element \(i\) in VRB. Word element \(i\) in VRT is set to all 1s if single-precision floating-point element i in VRA is greater than or equal to single-precision floating-point element i in VRB, and is set to all Os otherwise.

If the source element \(i\) in VRA or the source element \(i\) in VRB is a NaN, VRT is set to all Os, indicating "not greater than or equal to". If the source element in VRA and the source element i in VRB are both infinity with the same sign, VRT is set to all 1s, indicating "greater than or equal to".

\section*{Special Registers Altered:}

CR field 6
(if Rc=1)

\section*{Vector Compare Greater Than \\ Floating-Point VC-form}
\begin{tabular}{lll} 
vcmpgtfp & VRT,VRA,VRB & \((R c=0)\) \\
vcmpgtfp. & VRT,VRA,VRB & \((R c=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 04 & \[
6_{6} \text { VRT }
\] & \[
{ }_{11} \text { VRA }
\] & \[
{ }_{16} \text { VRB }
\] & RC & 710 & \\
\hline
\end{tabular}
```

do $\mathrm{i}=0$ to 127 by 32
$\operatorname{VRT}_{i: i+31} \leftarrow\left((\mathrm{VRA})_{i: i+31}>_{\text {fp }}(\mathrm{VRB})_{i: i+31}\right) ?{ }^{32} 1:{ }^{32} 0$
end
if $\mathrm{Rc}=1$ then do
$t \leftarrow\left(V R T={ }^{128} 1\right)$
$\mathrm{f} \leftarrow\left(\mathrm{VRT}={ }^{128} 0\right)$
$\mathrm{CR6} \leftarrow \mathrm{t}\|\mathrm{ObO}\| \mathrm{f} \| \mathrm{ObO}$
end

```

For each integer value i from 0 to 3 , do the following. Single-precision floating-point element i in VRA is compared to single-precision floating-point element i in VRB. Word element i in VRT is set to all 1s if single-precision floating-point element \(i\) in VRA is greater than single-precision floating-point element i in VRB, and is set to all Os otherwise.

If the source element \(i\) in VRA or the source element i in VRB is a NaN, VRT is set to all Os, indicating "not greater than". If the source element \(i\) in VRA and the source element \(i\) in VRB are both infinity with the same sign, VRT is set to all Os, indicating "not greater than".

Special Registers Altered:
CR field 6 . . . . . . . . . . . . . . . . . . . . . . . . . . (if Rc=1)

\subsection*{6.10.5 Vector Floating-Point Estimate Instructions}

\section*{Vector 2 Raised to the Exponent Estimate Floating-Point VX-form}
```

Vexptefp

| 4 | VRT,VRB |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | VRT |  |
| 11 |  | VRB |  | 394 |  |

$$
\begin{aligned}
& \text { do } i=0 \text { to } 127 \text { by } 32 \\
& \operatorname{VRT}_{i: i+31} \leftarrow \text { Power2EstimateSP }\left((\mathrm{VRB})_{i: i+31}\right) \\
& \text { end }
\end{aligned}
$$

```

For each integer value i from 0 to 3 , do the following.
The single-precision floating-point estimate of 2 raised to the power of single-precision floating-point element \(i\) in VRB is placed into word element \(i\) of VRT.

Let x be any single-precision floating-point input value. Unless \(x<-146\) or the single-precision floating-point result of computing 2 raised to the power \(x\) would be a zero, an infinity, or a QNaN, the estimate has a relative error in precision no greater than one part in 16. The most significant 12 bits of the estimate's significand are monotonic. An integral input value returns an integral value when the result is representable.

The result for various special cases of the source value is given below.
\begin{tabular}{cc} 
Value & Result \\
- Infinity & +0 \\
-0 & +1 \\
+0 & +1 \\
+ Infinity & + Infinity \\
NaN & QNaN
\end{tabular}

\section*{Special Registers Altered:}

None

\section*{Vector Log Base 2 Estimate \\ Floating-Point VX-form}


For each integer value i from 0 to 3 , do the following. The single-precision floating-point estimate of the base 2 logarithm of single-precision floating-point element \(i\) in VRB is placed into the corresponding word element of VRT.

Let x be any single-precision floating-point input value. Unless |x-1 | is less than or equal to 0.125 or the single-precision floating-point result of computing the base 2 logarithm of \(x\) would be an infinity or a QNaN , the estimate has an absolute error in precision (absolute value of the difference between the estimate and the infinitely precise value) no greater than \(2^{-5}\). Under the same conditions, the estimate has a relative error in precision no greater than one part in 8.

The most significant 12 bits of the estimate's significand are monotonic. The estimate is exact if \(\mathrm{x}=2^{\mathrm{y}}\), where y is an integer between -149 and +127 inclusive. Otherwise the value placed into the element of register VRT may vary between implementations, and between different executions on the same implementation.

The result for various special cases of the source value is given below.
\begin{tabular}{cc} 
Value & Result \\
- Infinity & QNaN \\
\(<0\) & QNaN \\
-0 & - Infinity \\
+0 & - Infinity \\
+ Infinity & +Infinity \\
NaN & QNaN
\end{tabular}

\section*{Special Registers Altered: \\ None}

\section*{Vector Reciprocal Estimate \\ Floating-Point VX-form}
vrefp VRT,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & & III & VRB & \multicolumn{2}{|c|}{266} & \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

do i=0 to 127 by 32
VRT
end

```

For each integer value i from 0 to 3 , do the following.
The single-precision floating-point estimate of the reciprocal of single-precision floating-point element i in VRB is placed into word element i of VRT.

Unless the single-precision floating-point result of computing the reciprocal of a value would be a zero, an infinity, or a QNaN, the estimate has a relative error in precision no greater than one part in 4096 .

Note that results may vary between implementations, and between different executions on the same implementation.

The result for various special cases of the source value is given below.
\begin{tabular}{cc} 
Value & Result \\
- Infinity & -0 \\
-0 & - Infinity \\
+0 & + Infinity \\
+ Infinity & +0 \\
NaN & QNaN
\end{tabular}

\section*{Special Registers Altered:}

None

\section*{Vector Reciprocal Square Root Estimate Floating-Point VX-form}
vrsqrtefp VRT,VRB
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & VRT & & III & VRB & \\
\hline 0 & & 630 & & & \\
\hline
\end{tabular}
do \(i=0\) to 127 by 32
\(\operatorname{VRT}_{i: i+31} \leftarrow\) RecipSquareRootEstimateSP \(\left((\mathrm{VRB})_{i: i+31}\right)\)
end
For each integer value i from 0 to 3 , do the following.
The single-precision floating-point estimate of the reciprocal of the square root of single-precision floating-point element \(i\) in VRB is placed into word element i of VRT .

Let \(x\) be any single-precision floating-point value. Unless the single-precision floating-point result of computing the reciprocal of the square root of \(x\) would be a zero, an infinity, or a QNaN, the estimate has a relative error in precision no greater than one part in 4096.

Note that results may vary between implementations, and between different executions on the same implementation.

The result for various special cases of the source value is given below.
\begin{tabular}{cc} 
Value & Result \\
- Infinity & QNaN \\
\(<0\) & QNaN \\
-0 & - Infinity \\
+0 & + Infinity \\
+ Infinity & +0 \\
NaN & QNaN
\end{tabular}

\section*{Special Registers Altered:}

None

\subsection*{6.11 Vector Exclusive-OR-based Instructions}

\subsection*{6.11.1 Vector AES Instructions}

This section describes a set of instructions that support the Federal Information Processing Standards Publica-

\section*{Vector AES Cipher VX-form}
vcipher
VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 04 & \[
{ }_{6} \text { VRT }
\] & \[
{ }_{11} \text { VRA }
\] & \[
{ }_{16} \text { VRB }
\] & 21 & 1288 \\
\hline State & \multicolumn{5}{|l|}{\(\leftarrow \mathrm{VR}\) [VRA]} \\
\hline RoundKey & \multicolumn{5}{|l|}{\(\leftarrow \mathrm{VR}[\mathrm{VRB}]\)} \\
\hline vtemp1 & \multicolumn{5}{|l|}{\(\leftarrow\) SubBytes(State)} \\
\hline vtemp2 & \multicolumn{5}{|l|}{\(\leftarrow\) ShiftRows(vtemp1)} \\
\hline vtemp3 & \multicolumn{5}{|l|}{\(\leftarrow\) MixColumns(vtemp2)} \\
\hline VR[VRT] & \multicolumn{5}{|l|}{\(\leftarrow\) vtemp3 ^ RoundKey} \\
\hline
\end{tabular}

Let State be the contents of VR[VRA], representing the intermediate state array during AES cipher operation.

Let Roundkey be the contents of VR[VRB], representing the round key.

One round of an AES cipher operation is performed on the intermediate State array, sequentially applying the transforms, SubBytes(), ShiftRows(), MixColumns(), and AddRoundKey () , as defined in FIPS-197.

The result is placed into VR[ VRT], representing the new intermediate state of the cipher operation.

\section*{Special Registers Altered:}

None
tion 197 Advanced Encryption Standard for encryption and decryption.

\section*{Vector AES Cipher Last VX-form}
vcipherlast VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & \({ }_{6}\) VRT & \({ }_{11}\) VRA & \({ }_{16}\) VRB & 21 & 1289 & 31 \\
\hline \multicolumn{7}{|l|}{State \(\leftarrow\) VR[VRA]} \\
\hline \multicolumn{7}{|l|}{RoundKey \(\leftarrow\) VR[VRB]} \\
\hline vtemp1 & \multicolumn{6}{|l|}{\(\leftarrow\) SubBytes(State)} \\
\hline vtemp2 & \multicolumn{6}{|l|}{\(\leftarrow\) ShiftRows(vtemp1)} \\
\hline VR[VRT] & \multicolumn{2}{|l|}{\(\leftarrow\) vtemp2 ^ RoundKey} & & & & \\
\hline
\end{tabular}

Let State be the contents of VR[VRA], representing the intermediate state array during AES cipher operation.

Let Round Key be the contents of VR[ VRB], representing the round key.

The final round in an AES cipher operation is performed on the intermediate State array, sequentially applying the transforms, SubBytes(), ShiftRows(), AddRoundKey(), as defined in FIPS-197.

The result is placed into VR[ VRT], representing the final state of the cipher operation.

\section*{Special Registers Altered:}

None

\section*{Vector AES Inverse Cipher VX-form}
vncipher VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \[
4
\] & \[
{ }_{6}{ }_{6} \text { VRT }
\] & \[
{ }_{11} \text { VRA }
\] & 16 & VRB & 21 & 1352 & 31 \\
\hline \multicolumn{8}{|l|}{State \(\leftarrow\) VR[VRA]} \\
\hline \multicolumn{8}{|l|}{RoundKey \(\leftarrow\) VR[VRB]} \\
\hline vtemp1 & \multicolumn{7}{|l|}{\(\leftarrow\) InvShiftRows(State)} \\
\hline vtemp2 & \multicolumn{7}{|l|}{\(\leftarrow\) InvSubBytes(vtemp1)} \\
\hline vtemp3 & \multicolumn{7}{|l|}{\(\leftarrow\) vtemp2 \(\wedge\) RoundKey} \\
\hline VR[VRT] & \multicolumn{7}{|l|}{\(\leftarrow\) InvMixColumns(vtemp3)} \\
\hline
\end{tabular}

Let State be the contents of VR[VRA], representing the intermediate state array during AES inverse cipher operation.

Let RoundKey be the contents of VR[ VRB], representing the round key.

One round of an AES inverse cipher operation is performed on the intermediate State array, sequentially applying the transforms, InvShiftRows(), InvSubBytes(), AddRoundKey(), and InvMixColumns(), as defined in FIPS-197.

The result is placed into VR[ VRT], representing the new intermediate state of the inverse cipher operation.

\section*{Special Registers Altered:}

None

\section*{Vector AES Inverse Cipher Last VX-form}
vncipherlast VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & \({ }_{6}\) VRT & \[
{ }_{11} \text { VRA }
\] & \({ }_{16}\) VRB & 21 & & 31 \\
\hline \multicolumn{7}{|l|}{State \(\leftarrow\) VR[VRA]} \\
\hline \multicolumn{7}{|l|}{RoundKey \(\leftarrow\) VR[VRB]} \\
\hline vtemp1 & \multicolumn{6}{|l|}{\(\leftarrow\) InvShiftRows(State)} \\
\hline vtemp2 & \multicolumn{6}{|l|}{\(\leftarrow\) InvSubBytes(vtemp1)} \\
\hline VR[VRT] & \multicolumn{6}{|l|}{\(\leftarrow\) vtemp2 ^ RoundKey} \\
\hline
\end{tabular}

Let State be the contents of VR[VRA], representing the intermediate state array during AES inverse cipher operation.

Let RoundKey be the contents of VR[VRB], representing the round key.

The final round in an AES inverse cipher operation is performed on the intermediate State array, sequentially applying the transforms, InvShiftRows(), InvSubBytes(), and AddRoundKey(), as defined in FIPS-197.

The result is placed into VR[ VRT] , representing the final state of the inverse cipher operation.

\section*{Special Registers Altered: \\ None}

\section*{Vector AES SubBytes VX-form}

\section*{vsbox VRT,VRA}


Let State be the contents of VR[ VRA], representing the intermediate state array during AES cipher operation.

The result of applying the transform, SubBytes() on St ate, as defined in FIPS-197, is placed into VR[ VRT].

\section*{Special Registers Altered: \\ None}

\subsection*{6.11.2 Vector SHA-256 and SHA-512 Sigma Instructions}

This section describes a set of instructions that support the Federal Information Processing Standards Publication 180-3 Secure Hash Standard.

\section*{Vector SHA-512 Sigma Doubleword VX-form}
vshasigmad
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & VRT, VRA,ST,SIX \\
\hline 0 & & 6 & VRA & ST & SIX & \\
16 & 1730 & & & \\
\hline
\end{tabular}
```

do i = 0 to 1
src \leftarrowVR[VRA].doubleword[i]
if ST=0 \& SIX.bit[2xi]=0 then // SHA-512 \sigma0 function
VR[VRT].dword[i] \leftarrow(src >>> 1) ^
(src>>> 8)^
(src >> 7)
if ST=0 \& SIX.bit[2xi]=1 then // SHA-512 \sigma1 function
VR[VRT].dword[i] \leftarrow (src >>> 19) ^
(src >>> 61) ^
(src >> 6)
if ST=1 \& SIX.bit[2xi]=0 then // SHA-512 \Sigma0 function
VR[VRT].dword[i] \leftarrow (src >>> 28)^
(src >>> 34)^
(src >>> 39)
if ST=1 \& SIX.bit[2xi]=1 then // SHA-512 \&1 function
VR[VRT].dword[i] \leftarrow(src >>> 14) ^
(src >>> 18)^
(src >>> 41)
end

```

For each integer value i from 0 to 1, do the following. When \(S T=0\) and bit \(2 x i\) of \(S I X\) is 0 , a SHA-512 \(\sigma 0\) function is performed on the contents of doubleword element \(i\) of VR[VRA] and the result is placed into doubleword element \(i\) of VR[ VRT].

When \(S T=0\) and bit \(2 x i\) of \(S I X\) is 1 , a SHA-512 \(\sigma 1\) function is performed on the contents of doubleword element \(i\) of VR[VRA] and the result is placed into doubleword element \(i\) of VR[ VRT].

When \(S T=1\) and bit \(2 x i\) of \(S I X\) is 0 , a SHA-512 \(\Sigma 0\) function is performed on the contents of doubleword element \(i\) of VR[VRA] and the result is placed into doubleword element \(i\) of VR[VRT].

When \(S T=1\) and bit \(2 \times i\) of \(S I X\) is 1 , a SHA- \(512 \Sigma 1\) function is performed on the contents of doubleword element \(i\) of VR[VRA] and the result is placed into doubleword element \(i\) of VR[ VRT].

Bits 1 and 3 of SIX are reserved.

\section*{Special Registers Altered:}

None

\section*{Vector SHA-256 Sigma Word VX-form}
vshasigmaw VRT,VRA,ST,SIX
\begin{tabular}{|l|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & ST & SIX & & 1666 \\
\hline 0 & & 6 & & 11 & & 16 & \\
\hline
\end{tabular}
do \(i=0\) to 3
\(\mathrm{src} \leftarrow \mathrm{VR}[\mathrm{VRA}]\). word[i]
if ST=0 \& SIX.bit[i]=0 then // SHA-256 \(\sigma 0\) function VR[VRT]. word[i] \(\leftarrow(\operatorname{src} \ggg 7) \wedge\)
\((\operatorname{src} \ggg 18) \wedge\)
(src >> 3)
if ST=0 \& SIX.bit[i]=1 then // SHA-256 o1 function
VR[VRT]. word[i] \(\leftarrow(\operatorname{src} \ggg 17) \wedge\)
\((\operatorname{src} \ggg 19)^{\wedge}\)
(src >> 10)
if ST=1 \& SIX.bit[i]=0 then // SHA-256 20 function VR[VRT]. word[i] \(\leftarrow(\operatorname{src} \ggg 2) \wedge\)
\((\operatorname{src} \ggg 13)^{\wedge}\)
(src >>> 22)
if ST=1 \& SIX.bit[i]=1 then // SHA-256 \(\Sigma 1\) function
VR[VRT]. word[i] \(\leftarrow(\operatorname{src} \ggg 6) \wedge\)
\((s r c \ggg 11) \wedge\)
(src >>> 25)

\section*{end}

For each integer value \(i\) from 0 to 3 , do the following. When \(S T=0\) and bit \(i\) of \(S I X\) is 0 , a SHA-256 \(\sigma 0\) function is performed on the contents of word element i of VR[VRA] and the result is placed into word element \(i\) of VR[ VRT].

When \(S T=0\) and bit \(i\) of \(S I X\) is 1 , a SHA-256 \(\sigma 1\) function is performed on the contents of word element \(i\) of VR[VRA] and the result is placed into word element \(i\) of VR[ VRT].

When \(S T=1\) and bit \(i\) of \(S I X\) is 0 , a SHA-256 \(\Sigma 0\) function is performed on the contents of word element \(i\) of VR[VRA] and the result is placed into word element \(i\) of \(\operatorname{VR[}\) VRT].

When \(S T=1\) and bit i of \(S I X\) is 1 , a SHA-256 \(\Sigma 1\) function is performed on the contents of word element \(i\) of VR[VRA] and the result is placed into word element \(i\) of VR[ VRT].

\section*{Special Registers Altered: \\ None}

\subsection*{6.11.3 Vector Binary Polynomial Multiplication Instructions}

This section describes a set of binary polynomial multi-ply-sum instructions. Corresponding elements are multiplied and the exclusive-OR of each even-odd pair of
products sum, useful for a variety of finite field arithmetic operations.

\section*{Vector Polynomial Multiply-Sum Byte VX-form}
```

vpmsumb VRT,VRA,VRB

| 4 | VRT | VRA | VRB | 1032 |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  |  |  |

```
```

if MSR.VEC=0 then Vector_Unavailable()

```
if MSR.VEC=0 then Vector_Unavailable()
do i = 0 to 15
do i = 0 to 15
    prod[i].bit[0:14] \leftarrow0
    prod[i].bit[0:14] \leftarrow0
    srcA \leftarrowVR[VRA].byte[i]
    srcA \leftarrowVR[VRA].byte[i]
    srcB }\leftarrow\textrm{VR[VRB].byte[i]
    srcB }\leftarrow\textrm{VR[VRB].byte[i]
    do j = 0 to 7
    do j = 0 to 7
        do k = 0 to j
        do k = 0 to j
            gbit \leftarrow srcA.bit[k] & srcB.bit[j-k]
            gbit \leftarrow srcA.bit[k] & srcB.bit[j-k]
            prod[i].bit[j] \leftarrow prod[i].bit[j] ^ gbit
            prod[i].bit[j] \leftarrow prod[i].bit[j] ^ gbit
            end
            end
    end
    end
    do j = 8 to 14
    do j = 8 to 14
        do k = j-7 to 7
        do k = j-7 to 7
            gbit \leftarrow (srcA.bit[k] & srcB.bit[j-k])
            gbit \leftarrow (srcA.bit[k] & srcB.bit[j-k])
            prod[i].bit[j] \leftarrow prod[i].bit[j] ^ gbit
            prod[i].bit[j] \leftarrow prod[i].bit[j] ^ gbit
            end
            end
    end
    end
end
end
do i = 0 to 7
do i = 0 to 7
    VR[VRT].hword[i] \leftarrow 0b0 |( (prod[2xi] ^ prod[2xi+1])
    VR[VRT].hword[i] \leftarrow 0b0 |( (prod[2xi] ^ prod[2xi+1])
end
```

end

```

For each integer value \(i\) from 0 to 15 , do the following. Let prod[i] be the 15-bit result of a binary polynomial multiplication of the contents of byte element i of VR[VRA] and the contents of byte element \(i\) of VR[ VRB].

For each integer value \(i\) from 0 to 7 , do the following. The exclusive-OR of \(\operatorname{prod}[2 x i]\) and \(\operatorname{prod}[2 x i+1]\) is placed in bits 1:15 of halfword element i of VR[ VRT]. Bit 0 of halfword element \(i\) of VR[VRT] is set to 0 .

\section*{Special Registers Altered:}

None

\section*{Vector Polynomial Multiply-Sum Doubleword VX-form}
```

vpmsumd VRT,VRA,VRB

```
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & & 1224 \\
\hline & & 6 & & 11 & & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
do i = 0 to 1
prod[i].bit[0:126] \leftarrow0
srcA \leftarrowVR[VRA].doubleword[i]
srcB \leftarrowVR[VRB].doubleword[i]
do j = 0 to 63
do k = 0 to j
gbit }\leftarrow\operatorname{srcA.bit[k] \& srcB.bit[j-k]
prod[i].bit[j] \leftarrow prod[i].bit[j] ^ gbit
end
end
do j = 64 to 126
do k = j-63 to 63
gbit \leftarrow (srcA.bit[k] \& srcB.bit[j-k])
prod[i].bit[j] \leftarrow prod[i].bit[j] ^ gbit
end
end
end
VR[VRT] \leftarrow 0b0 || (prod[0] ^ prod[1])

```

Let prod[0] be the 127-bit result of a binary polynomial multiplication of the contents of doubleword element 0 of VR[VRA] and the contents of doubleword element 0 of VR[ VRB].

Let prod[1] be the 127-bit result of a binary polynomial multiplication of the contents of doubleword element 1 of VR[VRA] and the contents of doubleword element 1 of VR[ VRB].

The exclusive-OR of \(\operatorname{prod}[0]\) and \(\operatorname{prod}[1]\) is placed in bits \(1: 127\) of VR[ VRT]. Bit 0 of VR[ VRT] is set to 0.

\section*{Special Registers Altered: \\ None}

\section*{Vector Polynomial Multiply-Sum Halfword VX-form}
vpmsumh VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1096} \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
```

do i = 0 to 7
prod.bit[0:30]}\leftarrow
srcA}\leftarrowVR[VRA], halfword[i]
srCB}\leftarrowVR[VRB], halfword[i]
do j=0 to 15
do k = 0 to j
gbit \leftarrow srcA.bit[k] \& srcB.bit[j-k]
prod[i].bit[j] \& prod[i].bit[j] ^ gbit
end
end
do j=16 to 30
do k= j-15 to 15
gbit \leftarrow(srcA.bit[k] \& srcB.bit[j-k])
prod[i].bit[j] \& prod[i].bit[j]^ gbit
end
end
end
VR[VRT].word[0] \& 0b0 ||(prod[0] ^ prod[1])
VR[VRT].word[1] \leftarrow 0b0 ||(prod[2] ^ prod[3])
VR[VRT].word[2] \leftarrow 0b0 ||(prod[4] ^ prod[5])
VR[VRT].word[3] \& 0b0 ||(prod[6] ^ prod[7])

```

For each integer value i from 0 to 7 , do the following. Let prod[i] be the 31-bit result of a binary polynomial multiplication of the contents of halfword element \(i\) of VR[VRA] and the contents of halfword element \(i\) of VR[VRB].

For each integer value i from 0 to 3 , do the following. The exclusive-OR of \(\operatorname{prod}[2 \times i]\) and \(\operatorname{prod}[2 x i+1]\) is placed in bits 1:31 of word element \(i\) of VR[VRT]. Bit 0 of word element \(i\) of \(V R[V R T]\) is set to 0 .

\section*{Special Registers Altered:}

None

\section*{Vector Polynomial Multiply-Sum Word VX-form}
vpmsumw VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|ccc|}
\hline 4 & VRT & VRA & VRB & \multicolumn{2}{|c|}{1160} & \\
\hline 0 & & 6 & 11 & & & 31 \\
\hline
\end{tabular}
do \(i=0\) to 3
\(\operatorname{prod}[i]\).bit[0:62] \(\leftarrow 0\)
\(\operatorname{srcA} \leftarrow \mathrm{VR}[\mathrm{VRA}]\).word[i]
\(\operatorname{srcB} \leftarrow \mathrm{VR}[\mathrm{VRB}] \cdot\) word[i]
do \(j=0\) to 31
do \(k=0\) to \(j\)
gbit \(\leftarrow \operatorname{srcA} . \operatorname{bit}[k] \& \operatorname{srcB} . \operatorname{bit}[j-k]\)
\(\operatorname{prod}[i] . \operatorname{bit}[j] \leftarrow \operatorname{prod}[i] . b i t[j] \wedge\) gbit
end
end
do \(j=32\) to 62
do \(k=j-31\) to 31 gbit \(\leftarrow(\operatorname{srcA} . \operatorname{bit}[k] \& \operatorname{srcB} . \operatorname{bit}[j-k])\)
\(\operatorname{prod}[\mathrm{i}] . \operatorname{bit}[j] \leftarrow \operatorname{prod}[\mathrm{i}] . \operatorname{bit}[j] \wedge\) gbit
end
end
end
VR[VRT].dword[0] \(\leftarrow 0 b 0 \|(\operatorname{prod}[0] \wedge \operatorname{prod}[1])\)
VR[VRT].dword[1] \(\leftarrow 0 \mathrm{bo} \|(\operatorname{prod}[2] \wedge \operatorname{prod}[3])\)
For each integer value i from 0 to 3 , do the following.
Let prod[i] be the 63-bit result of a binary polynomial multiplication of the contents of word element i of VR[VRA] and the contents of word element \(i\) of VR[VRB].

For each integer value i from 0 to 1 , do the following. The exclusive-OR of \(\operatorname{prod}[2 x i]\) and \(\operatorname{prod}[2 x i+1]\) is placed in bits 1:63 of doubleword element i of VR[VRT]. Bit 0 of doubleword element \(i\) of VR[VRT] is set to 0 .

\section*{Special Registers Altered:}

None

\subsection*{6.11.4 Vector Permute and Exclusive-OR Instruction}

\section*{Vector Permute and Exclusive-OR VA-form}
vpermxor VRT,VRA,VRB,VRC
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & VRC & & 45 \\
\hline 0 & & 6 & & & & \\
\hline
\end{tabular}
do \(i=0\) to 15
indexA \(\leftarrow \mathrm{VR}[\mathrm{VRC}]\).byte[i].bit[0:3]
index \(B \leftarrow V R[V R C]\).byte[i].bit[4:7]
src1 \(\leftarrow\) VR[VRA].byte[indexA]
src2 \(\leftarrow\) VR[VRB].byte[indexB]
VSR[VRT].byte[i] \(\leftarrow \operatorname{src1}\) ^ src2
end
For each integer value i from 0 to 15 , do the following. Let indexA be the contents of bits 0:3 of byte element \(i\) of VR[VRC].
Let indexB be the contents of bits \(4: 7\) of byte element \(i\) of VR[ VRC].

The exclusive OR of the contents of byte element indexA of VR[VRA] and the contents of byte element indexB of VR[VRB] is placed into byte element \(i\) of VR[ VRT].

Special Registers Altered:
None

\subsection*{6.12 Vector Gather Instruction}

\section*{Vector Gather Bits by Bytes by Doubleword VX-form}
vgbbd
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT, VRB \\
\hline 0 & & 6 & III & & VRB \\
16 & & 1292 & & \\
\hline
\end{tabular}
```

do i = 0 to 1
do j = 0 to 7
do k=0 to 7
b}\leftarrow\mathrm{ VSR[VRB].dword[i].byte[k].bit[j]
VSR[VRT].dword[i].byte[j].bit[k] \leftarrowb
end
end
end

```

Let \(\operatorname{src}\) be the contents of VR[VRB], composed of two doubleword elements numbered 0 and 1.

Let each doubleword element be composed of eight bytes numbered 0 through 7 .

An 8-bit \(\times 8\)-bit bit-matrix transpose is performed on the contents of each doubleword element of VR[VRB] (see Figure 107).

For each integer value i from 0 to 1 , do the following, The contents of bit 0 of each byte of doubleword element \(i\) of VR[ VRB] are concatenated and placed into byte 0 of doubleword element \(i\) of VR[VRT].

The contents of bit 1 of each byte of doubleword element \(i\) of VR[ VRB] are concatenated and placed into byte 1 of doubleword element i of VR[VRT].

The contents of bit 2 of each byte of doubleword element \(i\) of VR[VRB] are concatenated and placed into byte 2 of doubleword element i of VR[ VRT].

The contents of bit 3 of each byte of doubleword element \(i\) of VR[VRB] are concatenated and placed into byte 3 of doubleword element i of VR[ VRT].

The contents of bit 4 of each byte of doubleword element \(i\) of VR[VRB] are concatenated and placed into byte 4 of doubleword element i of VR[ VRT].

The contents of bit 5 of each byte of doubleword element \(i\) of VR[VRB] are concatenated and placed into byte 5 of doubleword element i of VR[ VRT].

The contents of bit 6 of each byte of doubleword element \(i\) of VR[VRB] are concatenated and placed into byte 6 of doubleword element \(i\) of VR[ VRT] .

The contents of bit 7 of each byte of doubleword element \(i\) of VR[VRB] are concatenated and placed into byte 7 of doubleword element i of VR[ VRT].

\section*{Special Registers Altered:}

None


Figure 107.Vector Gather Bits by Bytes by Doubleword

\subsection*{6.13 Vector Count Leading Zeros Instructions}


For each integer value i from 0 to 15 , do the following. A count of the number of consecutive zero bits starting at bit 0 of byte element i of VR[VRB] is placed into byte element i of VR[VRT]. This number ranges from 0 to 8 , inclusive.

\section*{Special Registers Altered:}

None

\section*{Vector Count Leading Zeros Halfword VX-form}

\section*{vclzh \\ VRT,VRB}
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & & III & VRB & & 1858 \\
\hline 0 & & 6 & & & & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
do i = 0 to 7
n}\leftarrow
do while n < 16
if VR[VRB].hword[i].bit[n] = 0b1 then leave
n}\leftarrow\textrm{n}+
end
VSR[VRT].hword[i] }\leftarrow\textrm{n
end

```

For each integer value \(i\) from 0 to 7 , do the following. A count of the number of consecutive zero bits starting at bit 0 of halfword element \(i\) of VR[VRB] is placed into halfword element i of VR[VRT]. This number ranges from 0 to 16 , inclusive.

\section*{Special Registers Altered:}

None

\section*{Vector Count Leading Zeros Word VX-form}

\author{
vclzw VRT,VRB
}
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & & III & VRB & \\
\hline 16 & 1922 & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
do i = 0 to 3
n}\leftarrow
do while n < 32
if VR[VRB].word[i].bit[n] = 0b1 then leave
n}\leftarrown+
end
VSR[VRT].word[i] \&n
end

```

For each integer value i from 0 to 3 , do the following.
A count of the number of consecutive zero bits starting at bit 0 of word element i of VR[VRB] is placed into word element i of VR[VRT]. This number ranges from 0 to 32 , inclusive.

\section*{Special Registers Altered:}

None

\section*{Vector Count Leading Zeros Doubleword VX-form}
VRIzd
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & \({ }_{6}\) VRT & & III & VRB & \\
\hline 0 & & 11 & & 1986 & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
do i = 0 to 1
n}\leftarrow
do while (n<64) \& (VR[VRB].dword[i].bit[n]=0b0)
n}\leftarrow\textrm{n}+
end
VSR[VRT].dword[i] }\leftarrow\textrm{n
end

```

For each integer value i from 0 to 1 , do the following. A count of the number of consecutive zero bits starting at bit 0 of doubleword element i of VR[VRB] is placed into doubleword element i of VR[VRT]. This number ranges from 0 to 64, inclusive.

Special Registers Altered:
None

\subsection*{6.14 Vector Count Trailing Zeros Instructions}
Vector Count Trailing Zeros Byte VX-form
Vctzb
\begin{tabular}{|l|l|l|l|l|ll|}
\hline 4 & VRT,VRB \\
0 & & 6 & 28 & & 11 & VRB \\
21 & 1538 & \\
\hline
\end{tabular}
if MSR. VEC=O then Vector_Unavailablel)
```

do i = O to 15
n}\leftarrow
do while n<8
if VR[VRB], byte[i],bit[7-n]=Ob1 then leave
n}\leftarrown+
end
VR[VRT], byte[i] \& Chop(EXTZ(n), 8)
end

```

For each integer value i from 0 to 15, do the following. A count of the number of consecutive zero bits starting at bit 7 of byte element i of VR[VRB] is placed into byte element i of VR[VRT]. This number ranges from 0 to 8 , inclusive.

\section*{Special Registers Altered:}

None

\section*{Vector Count Trailing Zeros Halfword VX-form}
vctzh VRT,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & & 29 & VRB & & 1538 \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
if MSR. VEC=O then Vector_Unavailablell
```

do i = 0 to 7
n}\leftarrow
do while n< 16
if VR[VRB].hword[i], bit[15-n]=0b1 then leave
n}\leftarrown+
end
VR[ VRT], hwor d[ i ] \&Chop( Extz(n), 16)
end

```

For each integer value i from 0 to 7 , do the following. A count of the number of consecutive zero bits starting at bit 15 of halfword element \(i\) of VR[VRB] is placed into halfword element i of VR[VRT]. This number ranges from 0 to 16 , inclusive.

\section*{Special Registers Altered:}

None

Vector Count Trailing Zeros Word
VX-form
VRT, VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & 30 & VRB & & 1538 \\
\hline 0 & & 6 & & 11 & \\
\hline
\end{tabular}
if MSR. VEC=O then Vector_Unavailablell
```

do i = 0 to 3
n}\leftarrow
do while n<32
if VR[VRB],word[i], bit[31-n] = Ob1 then I eave
n}\leftarrown+
end
VR[VRT].word[ i ] <Chop(EXTZ(n), 32)
end

```

For each integer value i from 0 to 3 , do the following. A count of the number of consecutive zero bits starting at bit 31 of word element i of VR[VRB] is placed into word element \(i\) of VR[ VRT]. This number ranges from 0 to 32 , inclusive.

Special Registers Altered:
None

\section*{Vector Count Trailing Zeros Doubleword VX-form}
VRT, VRB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 4 & VRT & & 31 & VRB & \\
\hline 0 & & 6 & 1538 & & 12 \\
\hline
\end{tabular}
```

if MSR. VEC=O then Vector_Unavailablell
do i = 0 to 1
n}\leftarrow
do while n<64
if VR[VNB],dword[i],bit[63-n] = Ob1 then leave
n}\leftarrown+
end
VR[ VRT].dword[i] \& Chop(EXTZ(n), 64)
end

```

For each integer value i from 0 to 1 , do the following. A count of the number of consecutive zero bits starting at bit 63 of doubleword element i of VR[VRB] is placed into doubleword element i of VR[VRT]. This number ranges from 0 to 64, inclusive.

\section*{Special Registers Altered:}

None

\subsection*{6.14.1 Vector Count Leading/Trailing Zero LSB Instructions}
```

Vector Count Leading Zero Least-Significant Bits Byte VX-form
vclzlsbb RT,VRB

```

```

if MSR.VEC=0 then Vector_Unavailable()

```
if MSR.VEC=0 then Vector_Unavailable()
count \leftarrow0
count \leftarrow0
do while count < 16
do while count < 16
    if (VR[VRB],byte[count],bit[7]=1) break
    if (VR[VRB],byte[count],bit[7]=1) break
    count \leftarrow count + 1
    count \leftarrow count + 1
end
end
GPR[RT] \leftarrow EXTZ64(count)
```

GPR[RT] \leftarrow EXTZ64(count)

```

Let count be the number of contiguous leading byte elements in VR[ VRB] having a zero least-significant bit.
count is placed into GPR[RT].
Special Registers Altered:
None

Vector Count Trailing Zero
Least-Significant Bits Byte VX-form
vctzlsbb RT,VRB

if MSR.VEC=0 then Vector _Unavailable()
count \(\leftarrow 0\)
do while count < 16
if (VR[VRB], byte[15-count], bit[7]=1) break
count \(\leftarrow\) count +1
end
GPR[RT] \(\leftarrow\) EXTZ64(count)

Let count be the number of contiguous trailing byte elements in VR[ VRB] having a zero least-significant bit.
count is placed into GPR[RT].
Special Registers Altered:
None

\subsection*{6.14.2 Vector Extract Element Instructions}
Vector Extract Unsigned Byte
Left-Indexed VX-form
vextublx
\begin{tabular}{|l|l|l|l|l|lll|}
\hline 4 & RT,RA,VRB \\
\hline 0 & 6 & & 11 & RA & 16 & VRB & \\
\hline 1 & 1549 & & 31 \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailable()
index \(\leftarrow\) GPR[RA], bit [60:63]
GPR[RT] \(\leftarrow\) EXTZ64 (VR[VRB], byte[index])
Let \(i\) index be the contents of bits 60:63 of GPR[RA].
The contents of byte element index of VR[VRB] are placed into bits \(56: 63\) of GPR[RT].

The contents of bits \(0: 55\) of GPR[RT] are set to 0 .
Special Registers Altered:
None

\section*{Vector Extract Unsigned Halfword Left-Indexed VX-form}
vextuhlx RT,RA,VRB
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 4 & RT & RA & VRB & & 1613 & \\
\hline 0 & & 6 & 11 & & 16 & \\
\hline
\end{tabular}
if MSR.VEC=0 then Vector_Unavailablel)
index \(\leftarrow\) GPR[RA], bit [60: 63]
GPR[RT] \(\leftarrow\) EXTZ64(VR[VRB], byte[index: index+1])
Let index be the contents of bits 60:63 of GPR[RA].
The contents of byte elements index:index+1 of VR[ VRB] are placed into bits \(48: 63\) of GPR[ RT].

The contents of bits \(0: 47\) of \(G P R[R T]\) are set to 0 .
If the value of \(i n d e x\) is greater than 14 , the results are undefined.

\section*{Special Registers Altered:}

None

\section*{Vector Extract Unsigned Byte Right-Indexed VX-form}
vextubrx RT,RA,VRB
\begin{tabular}{|l|l|l|c|c|ccc|}
\hline 4 & \multicolumn{1}{|c|}{ RT } & \multicolumn{2}{c|}{ RA } & \({ }_{11}\) VRB & & 1805 & \\
\hline 0 & & 6 & & & 21 & & 31 \\
\hline
\end{tabular}
if MSR.VEC=0 then Vector _Unavailablell
index \(\leftarrow\) GPR[RA], bit [60:63]
GPR[RT] \(\leftarrow\) EXTZ64(VR[VRB].byte[15-index])
Let index be the contents of bits 60:63 of GPR[RA].
The contents of byte element 15 -index of VR[VRB] are placed into bits 56:63 of GPR[RT].

The contents of bits \(0: 55\) of GPR[RT] are set to 0 .
Special Registers Altered:
None

\section*{Vector Extract Unsigned Halfword Right-Indexed VX-form}
vextuhrx RT,RA,VRB
\begin{tabular}{|l|l|l|l|l|lll|}
\hline 4 & RT & RA & VRB & & 1869 & \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
if MSR.VEC=O then Vector Unavailablel)
index \(\leftarrow\) GPR[RA], bit [60:63]
GPR[RT] \(\leftarrow\) EXTZ64(VR[VRB].byte[14-index:15•index])
Let \(i n d e x\) be the contents of bits 60:63 of GPR[RA].
The contents of byte elements 14 -index:15-index of VR[ VRB] are placed into bits 48:63 of GPR[ RT].

The contents of bits \(0: 47\) of GPR[ RT] are set to 0 .
If the value of index is greater than 14, the results are undefined.

\section*{Special Registers Altered:}

None

\section*{Vector Extract Unsigned Word Left-Indexed VX-form}
vextuwlx RT,RA,VRB
\begin{tabular}{|l|l|l|l|l|l|}
\hline 4 & 6 & RT & RA & VRB & \\
\hline 0 & & 1677 & \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailablel)
index \(\leftarrow\) GPR[ RA], bit [60: 63]
GPR[RT] EXTZ64(VR[VRB]. byte[index: index+3])
Let i ndex be the contents of bits 60:63 of GPR[ RA] .
The contents of byte elements index:index+3 of VR[ VRB] are placed into bits 32:63 of GPR[ RT].

The contents of bits \(0: 31\) of GPR[RT] are set to 0 .
If the value of i index is greater than 12, the results are undefined.

\section*{Special Registers Altered:}

None

\section*{Vector Extract Unsigned Word Right-Indexed VX-form}
vextuwrx RT,RA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ RT } & \multicolumn{1}{c|}{ RA } & \multicolumn{1}{|c|}{ VRB } & \multicolumn{2}{|c|}{1933} & 31 \\
\hline
\end{tabular}
if MSR.VEC=0 then Vector_Unavailablel)
index \(\leftarrow\) GPR[RA], bit [60:63]
GPR[RT] \(\leftarrow\) EXTZ64(VR[VRB]. byte[12-index: 15-index])
Let index be the contents of bits 60:63 of GPR[ RA].
The contents of byte elements index:index+3 of VR[ VRB] are placed into bits 32:63 of GPR[ RT].

The contents of bits \(0: 31\) of GPR[RT] are set to 0 .
If the value of index is greater than 12, the results are undefined.

Special Registers Altered:
None

\subsection*{6.15 Vector Population Count Instructions}

\section*{Vector Population Count Byte VX-form}
vpopentb VRT,VRB
\begin{tabular}{|c|c|c|c|cc|}
\hline 4 & VRT & & III & VRB & \\
\hline 0 & & 1795 & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
do i = 0 to 15
n}\leftarrow
do j = 0 to 7
n}\leftarrow\textrm{n}+\textrm{VR[VRB].byte[i].bit[j]
end
VSR[VRT].byte[i] \&n
end

```

For each integer value i from 0 to 15 , do the following. A count of the number of bits set to 1 in byte element i of \(V R[V R B]\) is placed into byte element i of VR[VRT]. This number ranges from 0 to 8 , inclusive.

\section*{Special Registers Altered:}

None

\section*{Vector Population Count Doubleword VX-form}
vpopentd VRT,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & & III & VRB & & 1987 \\
\hline 0 & & 6 & & & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
do i = 0 to 1
n}\leftarrow
do j = 0 to 63
n}\leftarrow\textrm{n}+\mathrm{ VR[VRB].dword[i].bit[j]
end
VSR[VRT].dword[i] < n
end

```

For each integer value i from 0 to 1 , do the following.
A count of the number of bits set to 1 in doubleword element i of VR[VRB] is placed into doubleword element i of VR[VRT]. This number ranges from 0 to 64 , inclusive.

\section*{Special Registers Altered:}

None

\section*{Vector Population Count Halfword VX-form}
```

vpopenth VRT,VRB

| 4 | VRT |  | IIII | VRB |  | 1859 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 61 |  |  |  |  |

if MSR.VEC=0 then Vector_Unavailable()
do i = 0 to 7
n}\leftarrow
do j = 0 to 15
n}\leftarrow\textrm{n}+\operatorname{VR[VRB].hword[i].bit[j]
end
VSR[VRT].hword[i] \leftarrown
end

```

For each integer value i from 0 to 7, do the following.
A count of the number of bits set to 1 in halfword element \(i\) of VR[VRB] is placed into halfword element \(i\) of VR[ VRT]. This number ranges from 0 to 16 , inclusive.

\section*{Special Registers Altered:}

None

\section*{Vector Population Count Word VX-form}
vpopentw VRT,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \({ }_{6}\) VRT & & III & & VRB & \\
\hline 16 & 1923 & \\
\hline
\end{tabular}
```

if MSR.VEC=0 then Vector_Unavailable()
do i = 0 to 3
n}\leftarrow
do j=0 to 31
n}\leftarrow\textrm{n}+\textrm{VR[VRB].word[i].bit[j]
end
VSR[VRT].word[i] \leftarrow n
end

```

For each integer value ifrom 0 to 3 , do the following. A count of the number of bits set to 1 in word element \(i\) of VR[VRB] is placed into word element \(i\) of VR[VRT]. This number ranges from 0 to 32, inclusive.

\section*{Special Registers Altered:}

None

\subsection*{6.16 Vector Bit Permute Instruction}
```

Vector Bit Permute Doubleword VX-form
vbpermd VRT,VRA,VRB

| 4 | VRT |  | VRA | VRB | 1484 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  |  | 11 |  |  |  |

if MSR. VEC=O then Vector_Unavailable()
do i = 0 to l
do j = 0 to 7
index \leftarrowVR[ VRB], dword[ i], byte[j]
if index < 64 then
perm.bit[j] \& VR[VRA].dword[i],bit[index]
else
permbit[j]}\leftarrow
end
VR[VRT].dword[i] \& EXTZ64(perm)
end

```

For each integer value i from 0 to 1 , and for each integer value \(j\) from 0 to 7 , do the following.

Let index be the contents of byte sub-element \(j\) of doubleword element \(i\) of VR[VRB].

If index is less than 64, then the contents of bit index of doubleword i of VR[VRA] are placed into bit \(56+j\) of doubleword element i of VR[VRT]. Otherwise, bit \(56+j\) of doubleword element \(i\) of VR[VRT] is set to 0.

The contents of bits 0:55 of doubleword element i of VR[VRT] are set to 0 .

\section*{Special Registers Altered:}

None

Vector Bit Permute Quadword VX-form
vbpermq VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & VRB & \multicolumn{2}{|c|}{1356} \\
\hline 0 & & 6 & & & & \\
\hline
\end{tabular}
```

if MSR.VEC=O then Vector_Unavailable()
do i = 0 to 15
index \leftarrow VR[VRB],byte[i]
if index < 128 then
perm,bit[i] \&VR[VRA],bit[index]
else
perm.bit[i]}\leftarrow
end
VR[VRT].dword[0] \leftarrowChop(EXTZ(perm),64)
VR[VRT].dword[1]}\leftarrow0\times0000_0000_0000_000

```

For each integer value \(i\) from 0 to 15 , do the following.
Let index be the contents of byte element i of VR[VRB].

If index is less than 128 , then the contents of bit index of VR[VRA] are placed into bit \(48+i\) of doubleword element i of VR[VRT]. Otherwise, bit \(48+\mathrm{i}\) of doubleword element \(i\) of VR[VRT] is set to 0 .

The contents of bits \(0: 47\) of VR[ VRT] are set to 0 .
The contents of bits \(64: 127\) of VR[ VRT] are set to 0.

\section*{Special Registers Altered:}

None

\section*{Programming Note}

The fact that the permuted bit is 0 if the corresponding index value exceeds 127 permits the permuted bits to be selected from a 256 -bit quantity, using a single index register. For example, assume that the 256 -bit quantity \(Q\), from which the permuted bits are to be selected, is in registers v2 (high-order 128 bits of \(Q\) ) and \(v 3\) (low-order 128 bits of \(Q\) ), that the index values are in register \(v 1\), with each byte of \(v 1\) containing a value in the range \(0: 255\), and that each byte of register v4 contains the value 128. The following code sequence selects eight permuted bits from \(Q\) and places them into the low-order byte of \(v 6\).
```

vbpermq v6,v1,v2 \# select fromhigh-order half
of Q
vxor vo,v1,v4 \# adjust index values
vbpermg v5,v0,v3 \# select fromlow-order half
of Q
vor v6,v6,v5 \# merge the two selections

```

\subsection*{6.17 Decimal Integer Instructions}

A valid encoding of a packed decimal integer value requires the following properties.
- Each of the 31 4-bit digits of the operand's magnitude (bits 0:123) must be in the range 0-9.
- The sign code (bits 124:127) must be in the range 10-15.

Source operands with sign codes of 0b1010, Ob1100, \(0 b 1110\), and \(0 b 1111\) are interpreted as positive values.

Source operands with sign codes of 0b1011 and 0b1101 are interpreted as negative values.

Positive and zero results are encoded with a either sign code of \(0 b 1100\) or Ob1111, depending on the preferred sign (indicated as an immediate operand).

Negative results are encoded with a sign code of Ob1101.

\subsection*{6.17.1 Decimal Integer Arithmetic Instructions}

The Decimal Integer Arithmetic instructions operate on decimal integer values only in signed packed decimal format. Signed packed decimal format consists of 31 4-bit base-10 digits of magnitude and a trailing 4-bit
sign code. Operations are performed as sign-magnitude, and produce a decimal result placed in a Vector Register (i.e., bcdadd, bcdsub).

\section*{Decimal Add Modulo VX-form}
bcdadd. VRT,VRA,VRB,PS
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & VRT & VRA & VRB & 1 & 1 & PS \\
21 & & 1 & 1 & \\
\hline
\end{tabular}
if MSR. VEC=0 then Vector_Unavailablel)
VR[ VRT] \(\leftarrow\) Si gned_BCD_Add (VR[ VRA], VR[ VRB], PS)
CR.bit[56] \(\leftarrow\) inv_flag? ObO: It_flag
CR.bit[57] \(\leftarrow\) inv_flag? ObO: gt_flag
CR.bit[58] \(\leftarrow\) inv_flag? ObO: eq_flag
CR.bit[59] \(\leftarrow 0 \times\) - \(\dagger\) lag \(\mid\) inv_flag
Let srcl be the decimal integer value in VR[VRA].
Let \(s r^{\prime} 2\) be the decimal integer value in VR[VRB].
srCl is added to \(\mathrm{src2}\).
If the unbounded result is equal to zero, do the following.

If \(P S=0\), the sign code of the result is set to \(0 b 1100\).
If \(P S=1\), the sign code of the result is set to \(0 b 1111\).
CR field 6 is set to \(0 b 0010\).
If the unbounded result is greater than zero, do the following.

If \(P S=0\), the sign code of the result is set to \(0 b 1100\).
If \(P S=1\), the sign code of the result is set to \(0 b 1111\).
If the operation overflows, \(C R\) field 6 is set to \(0 b 0101\). Otherwise, \(C R\) field 6 is set to \(0 b 0100\).

If the unbounded result is less than zero, do the following.

The sign code of the result is set to 0 b1101.
If the operation overflows, \(C R\) field 6 is set to \(0 b 1001\). Otherwise, CR field 6 is set to \(0 b 1000\).

The low-order 31 digits of the magnitude of the result are placed in bits \(0: 123\) of VR[ VRT].

The sign code is placed in bits 124:127 of VR[ VRT].
If either srcl orsrc2 is an invalid encoding of a 31 -digit signed decimal value, the result is undefined and CR field 6 is set to \(0 b 0001\).

\section*{Special Registers Altered:}

CR field 6

\section*{Decimal Subtract Modulo VX-form}
bcdsub. VRT,VRA,VRB,PS
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 4 & \({ }_{6}\) VRT & \({ }_{11}\) VRA & \({ }_{16}\) VRB & \[
\left|\begin{array}{l|l|}
\hline 1 & \mid P S \\
2122
\end{array}\right|
\] & 65 & \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailablell
VR[ VRT] \(\leftarrow\) Signed_ BCD_Subtract (VR[VRA], VR[ VRB], PS)
CR.bit[56] \(\leftarrow\) inv_flag? Obo: It_flag
CR.bit[57] \(\leftarrow\) inv_flag? Obo: gt_flag
CR.bit[58] \(\leftarrow\) inv_flag? ObO: eq_flag
CR. bit[59] \(\leftarrow 0 x_{\text {_ }} \mathrm{f}|\mathrm{ag}|\) inv_flag
Let srcl be the decimal integer value in VR[VRA].
Let \(\operatorname{src} 2\) be the decimal integer value in VR[ VRB].
srCl is subtracted by src .
If the unbounded result is equal to zero, do the following.

If \(P S=0\), the sign code of the result is set to \(0 b 1100\).
If \(P S=1\), the sign code of the result is set to \(0 b 1111\).
\(C R\) field 6 is set to \(0 b 0010\).
If the unbounded result is greater than zero, do the following.

If \(P S=0\), the sign code of the result is set to \(0 b 1100\).
If \(P S=1\), the sign code of the result is set to \(0 b 1111\).
If the operation overflows, \(C R\) field 6 is set to \(0 b 0101\). Otherwise, CR field 6 is set to \(0 b 0100\).

If the unbounded result is less than zero, do the following.

The sign code of the result is set to 0 b1101.
If the operation overflows, \(C R\) field 6 is set to \(0 b 1001\). Otherwise, CR field 6 is set to \(0 b 1000\).

The low-order 31 digits of the magnitude of the result are placed in bits \(0: 123\) of VR[ VRT].

The sign code is placed in bits 124:127 of VR[ VRT].
If either \(\operatorname{srcl}\) or \(\mathrm{SrCl}_{2}\) is an invalid encoding of a 31 -digit signed decimal value, the result is undefined and CR field 6 is set to \(0 b 0001\).

\section*{Special Registers Altered:}

CR field 6

\subsection*{6.17.2 Decimal Integer Format Conversion Instructions}

if MSR.VEC=0 then Vector Unavailablell
srcsign \(\leftarrow(V R[V R B]\).hword[7] \(=0 \times 0020)\)
\(e q_{-} \mathrm{f} \operatorname{lag} \leftarrow 1\)
1* check for valid sign */
inv_flag \(\leftarrow(\) VR [VRB].hword[7] ! \(=0 \times 002 B) \&\)
(VR[ VRB], hword[7]!=0×002D)
do \(\mathrm{i}=0\) to 6
eq_flag \(\leftarrow\) eq_flag \& (VR[VRB].hword[i] \(=0 \times 0030\) )
|* check for valid digit *।
inv_flag \&inv_flag | (VR[VRB].hword[i]<0x0030)
| (VR[VRB].hword[i] >0x0039)
end

It_flag \(\leftarrow\left(e q_{2} f a g=0\right) \&\left(s r c \_\right.\)sign=1)
gt_flag \(\leftarrow\left(e q_{-}^{-} f a g=0\right) \&\left(\operatorname{sic}_{2} \operatorname{sign}=0\right)\)
do \(\mathrm{i}=0\) to 23
result, nibble \([i] \leftarrow 0 \times 0\)
end
do \(\mathrm{i}=0\) to 6
result, nibble[i +24] \(\leftarrow\) VR[VRB], hword[i], nibble[3]
end
result. nibble[31] \(\leftarrow(\) src_sign=0) ? \(((P S=0) ? O \times C: O X F): O \times D\)
VR[VRT] \(\leftarrow\) inv_flag? undefined: result
CR.bit[56] \(\leftarrow\) inv_flag? Obo: It_flag
CR.bit[57] \(\leftarrow\) inv_flag? ObO: gt_flag
CR.bit[58] \(\leftarrow\) inv_flag? ObO: eq_flag
CR.bit[59] \(\leftarrow\) inv_flag

Let \(s r^{c}\) be the national decimal value in VR[VRB].
\(s r c\) is placed in VR[VRT] in packed decimal format.
A valid encoding of a national decimal value requires the following.
- The contents of halfword 7 (sign code) must be either \(0 \times 002 \mathrm{~B}\) or \(0 \times 002 \mathrm{D}\).
- The contents of halfwords 0 to 6 must be in the range \(0 \times 0030\) to \(0 \times 0039\).

National decimal values having a sign code of \(0 \times 002 \mathrm{~B}\) are interpreted as positive values.

National decimal values having a sign code of \(0 \times 0020\) are interpreted as negative values.

For each integer value i from 0 to 23 , do the following. The contents of nibble element \(i\) of VR[VRT] are set to \(0 \times 0\).

For each integer value i from 0 to 6 , do the following. The contents of nibble 3 of halfword element i of \(\operatorname{src}\) are placed into nibble element \(i+24\) of VR[ VRT].

For \(P S=0\), the contents of nibble element 31 (i.e., sign code) of VR[VRT] are set to \(0 \times C\) for positive values and to \(0 \times D\) for negative values.

For \(P S=1\), the contents of nibble element 31 (i.e., sign code) of VR[VRT] are set to \(0 \times F\) for positive values and to \(0 \times D\) for negative values.

CR field 6 is set to reflect \(\operatorname{src}\) compared to zero.
If \(\operatorname{src}\) is an invalid encoding of a national decimal value, the contents of VR[VRT] are undefined and \(C R\) field 6 is set to \(0 b 0001\).

Special Registers Altered:
CR field 6

\section*{Decimal Convert From Zoned VX-form bcdcfz. VRT,VRB,PS \\ \begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & \multicolumn{2}{|c|}{6} & \multicolumn{2}{|c|}{ VRB } & 1 \\
\hline 0 & & PS & & 385 & \\
\hline
\end{tabular}}
```

if MSR.VEC=O then Vector_Unavailable()
|* check for valid sign *I
inv_flag \& ((VR[VRB], byte[15], ni bble[0] < OxA) \& (PS=1)) |
(VR[VRB], byte[15], ni bble[1] > Ox9)
|* check for valid digits */
MIN}\leftarrow(PS=0) ? OX30: OXF
MAX \leftarrow(PS=0) ? OX30: OXF9
do i = 0 to 14
inv_flag<inv_flag|(VR[VRB],byte[i]<MMN)
| (VR[VRB], byte[i] > MAX)
end
if PS=0 then
src_sign}\leftarrow VR[VRB],nibble[30],bit[1
else
src_sign \& (VR[VRB],nibble[30]=Ob1011)
(VR[VRB], nibbl e[30] = 0b1101)

```
eq_fag \(\leftarrow 1\)
do \(i=0\) to 14
    result.nibble[i] \(\leftarrow 0 \times 0\)
end
do \(i=0\) to 15
    result. nibble[i +15\(] \leftarrow\) VR[VRB]. byte[i], nibble[1]
    eq_flag \(\leftarrow e q_{-} f l a g \&(V R[V R B]\), byte[i], nibble[1]=0×0)
end
It_fag \(\leftarrow\left(e q_{-} f(a g=0) \&\left(s r c \_i g n=1\right)\right.\)
gt_flag \(\leftarrow\left(e q_{-}^{-} f(a g=0) \&\left(s c_{1}^{-} \operatorname{sign}=0\right)\right.\)
result. nibble[31] \(\leftarrow(\) src_sign=0 \()\) ? OXC : OXD
VR[VRT] \(\leftarrow\) inv_flag? undefined: result
CR. bit[56] \(\leftarrow\) inv_flag? ObO: It_flag
CR, bit[57] \(\leftarrow\) inv_flag? ObO: gt_flag
CR.bit[58] \(\leftarrow\) inv_flag? ObO: eq_flag
CR.bit[59] \(\leftarrow\) inv_flag

Let \(\operatorname{src}\) be the zoned decimal value in VR[ VRB].
\(s r c\) is placed in VR[ VRT] in packed decimal format.
When \(P S=0\), do the following.
A valid encoding of a zoned decimal value requires the following.
- The contents of bits 0:3 of byte 15 (sign code) can be any value in the range \(0 \times 0\) to \(0 \times F\).
- The contents of bits \(0: 3\) of bytes 0 to 14 must be the value \(0 \times 3\).
- The contents of bits \(4: 7\) of bytes 0 to 15 must be a value in the range \(0 \times 0\) to \(0 \times 9\).

Zoned decimal values having a sign code of \(0 \times 0\), \(0 \times 1,0 \times 2,0 \times 3,0 \times 8,0 \times 9,0 \times A\), or \(0 \times B\) are interpreted as positive values.

Zoned decimal values having a sign code of \(0 \times 4\), \(0 \times 5,0 \times 6,0 \times 7,0 \times C, 0 \times D, 0 \times E\), or \(0 \times F\) are interpreted as negative values.

When \(P S=1\), do the following.
A valid encoding of a zoned decimal source operand requires the following.
- The contents of bits 0:3 of byte 15 (sign code) must be a value in the range \(0 \times A\) to \(0 \times F\).
- The contents of bits \(0: 3\) of bytes 0 to 14 must be the value \(0 \times F\).
- The contents of bits \(4: 7\) of bytes 0 to 15 must be a value in the range \(0 \times 0\) to \(0 \times 9\).

Zoned decimal source operands having a sign code of \(0 \times A, O \times C, O \times E\), or \(0 \times F\) are interpreted as positive values.

Zoned decimal source operands having a sign code of \(0 \times B\) or \(0 \times D\) are interpreted as negative values.

Positive packed decimal results are returned with a sign code of \(0 \times C\).

Negative packed decimal results are returned with a sign code of \(0 \times D\).

For each integer value i from 0 to 14 ,
The contents of nibble element i of VR[VRT] are set to \(0 \times 0\).

For each integer value i from 0 to 15 ,
The contents of nibble 1 of byte element \(i\) of src are placed into nibble element \(i+15\) of VR[ VRT].

CR field 6 is set to reflect src compared to zero.
If src is an invalid encoding of a zoned decimal value, the contents of VR[VRT] are undefined and CR field 6 is set to 0 b 0001 .

\section*{Special Registers Altered: \\ CR field 6}

\section*{Decimal Convert To National VX-form}
bcdctn. \({ }^{\text {VRT,VRB }}\)
\begin{tabular}{|l|l|l|l|l|l|l|lll|}
\hline 4 & VRT & & 5 & VRB & 1 & \(/\) & & 385 & \\
\hline 0 & & 6 & & 11 & 21 & 22 & 23 & & 31 \\
\hline
\end{tabular}
if MSR.VEC=0 then Vector_Unavailable()
\(0 x \_f a g \leftarrow 0\)
do \(\mathrm{i}=0\) to 23
ox_flag \(\leftarrow 0 x_{\text {_ }}\) flag \(\mid(\) VR[VRB].nibble[i]! \(=0 \times 0)\)
end
inv_flag \(\leftarrow(\) VR[VRB], nibble[31]<OXA)
do \(\bar{i}=0\) to 30
inv_flag \(\leftarrow\) inv_flag \(\mid\) (VR[VRB].nibble[i] >Oxg)
end
srcssign \(\leftarrow(\) VR[VRB].nibble[31] \(=0 \times B) \mid\)
(VR[VRB].nibble[31] \(=0 \times D\) )
eq_flag \(\leftarrow(\) VR[VRB]. nibble[0:30] \(=0)\)
|t_flag \(\leftarrow\left(e q_{-} f a g=0\right) \&\left(s r c_{\_}\right.\)sign=1)
\(g t_{\_}^{-} f a g \leftarrow\left(e q_{1}^{-} f a g=0\right) \&\left(\operatorname{sic}_{1}^{-} \operatorname{sign}=0\right)\)
do \(\mathrm{i}=0\) to 6
result.hword[i].nibble[0:2] 0x003
result, hword[i].nibble[3] \(\leftarrow\) VR[VRB], nibble[i+24]
end
result.hword[7] \(\leftarrow(\) src_sign \(=1) ? 0 \times 0020: ~ O \times 002 B\)
VR[VRT] \(\leftarrow\) inv_flag ? undefined: result
CR.bit[56] \(\leftarrow\) inv_flag? ObO: It_flag
CR.bit[57] \(\leftarrow\) inv_flag? ObO: gt_flag
CR.bit[58] \(\leftarrow\) inv_flag ? ObO: eq_flag
CR. bit[59] \(\leftarrow\) inv_flag \(\mid\) ox_flag
Let \(\operatorname{src}\) be the packed decimal value in VR[VRB].
\(s r c\) is placed into VR[VRT] in national decimal format.

A valid encoding of a signed packed decimal value requires the following.
- The contents of nibble 31 (sign code) must be a value in the range \(0 \times A\) to \(0 \times F\).
- The contents of each nibble 0-30 must be a value in the range \(0 \times 0\) to \(0 \times 9\).

Packed decimal values with sign codes of \(0 \times A, O \times C\), 0 XE , or 0 XF are interpreted as positive values.

Packed decimal values with sign codes of \(0 \times B\) or \(0 \times D\) are interpreted as negative values.

Values greater in magnitude than \(10^{7}-1\) are too large to be represented in national decimal format.

For each integer value i from 0 to 6, do the following. The value \(0 \times 003\) is placed into nibbles \(0: 2\) of halfword element \(i\) of VR[VRT].

The contents of nibble element \(i+24\) of \(V R[\) VRB] are placed into nibble 3 of halfword element i of VR[ VRT].

The contents of halfword element 7 (i.e., sign code) of VR[VRT] are set to \(0 \times 002 \mathrm{~B}\) for positive values and to \(0 \times 002 \mathrm{D}\) for negative values.

CR field 6 is set to reflect src compared to zero, including whether or not src is too large to be represented in national decimal format.

If src is an invalid encoding of a packed decimal value, the contents of VR[ VRT] are undefined and CR field 6 is set to 0b0001.

Special Registers Altered:
CR field 6
```

Decimal Convert To Zoned VX-form
bcdctz. VRT,VRB,PS

| 4 | VRT |  | 4 |  | VRB | 1 | PS |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 41 |  | 385 |  |  |  |  |

```
```

if MSR.VEC=0 then Vector_Unavailable()

```
if MSR.VEC=0 then Vector_Unavailable()
inv_flag \leftarrow(VR[VRB], nibble[31] < OxA)
inv_flag \leftarrow(VR[VRB], nibble[31] < OxA)
do i = 0 to 30
do i = 0 to 30
    inv_flag & inv_flag | (VR[VRB].nibble[i] > Oxg)
    inv_flag & inv_flag | (VR[VRB].nibble[i] > Oxg)
end
end
0x_flag<0
0x_flag<0
do i = 0 to 15
do i = 0 to 15
    Ox_flag<ox_flag | (VR[VRB],nibble[i]!=OxO)
    Ox_flag<ox_flag | (VR[VRB],nibble[i]!=OxO)
end
end
src_sign & (VR[VRB], nibble[31]= OxB) |
src_sign & (VR[VRB], nibble[31]= OxB) |
    (VR[VRB],nibble[31] = OxD)
    (VR[VRB],nibble[31] = OxD)
eq_flag \leftarrow(VR[VRB],nibble[0:30]=0)
eq_flag \leftarrow(VR[VRB],nibble[0:30]=0)
|t_flag \leftarrow(eq_flag=0) & (src_sign=1)
|t_flag \leftarrow(eq_flag=0) & (src_sign=1)
gt_flag}\leftarrow(e\mp@subsup{q}{_}{-}f(ag=0)&(src_sign=0
gt_flag}\leftarrow(e\mp@subsup{q}{_}{-}f(ag=0)&(src_sign=0
do i = 0 to 14
do i = 0 to 14
    result.byte[i], nibble[0]}\leftarrow(PS=0) ? Ox3: OXF
    result.byte[i], nibble[0]}\leftarrow(PS=0) ? Ox3: OXF
    result, byte[i], nibble[1] &VR[VRB], nibble[i +15]
    result, byte[i], nibble[1] &VR[VRB], nibble[i +15]
end
end
if src.sign=0 then
if src.sign=0 then
    result.byte[15], nibble[0] \leftarrow(PS=0) ? OX3 : OxC
    result.byte[15], nibble[0] \leftarrow(PS=0) ? OX3 : OxC
else
else
    result.byte[15], nibble[0]}\leftarrow(PS=0) ? OX7: OxD
    result.byte[15], nibble[0]}\leftarrow(PS=0) ? OX7: OxD
result, byte[15], nibble[1] \leftarrow VR[VRB], nibble[30]
result, byte[15], nibble[1] \leftarrow VR[VRB], nibble[30]
VR[VRT] \leftarrowinv_flag? undefined: result
VR[VRT] \leftarrowinv_flag? undefined: result
CR,bit[56] \leftarrow inv_flag ? ObO: It_flag
CR,bit[56] \leftarrow inv_flag ? ObO: It_flag
CR.bit[57] \leftarrowinv_flag? ObO: gt flag
CR.bit[57] \leftarrowinv_flag? ObO: gt flag
CR.bit[58] \leftarrow inv_flag? Obo: eq_flag
CR.bit[58] \leftarrow inv_flag? Obo: eq_flag
CR.bit[59]}\leftarrow\mathrm{ inv_flag | ox_flag
```

CR.bit[59]}\leftarrow\mathrm{ inv_flag | ox_flag

```

Let \(s i c\) be the packed decimal value in VR[VRB].
src is placed into VR[VRT] in zoned decimal format.
A valid encoding of a signed packed decimal value requires the following.
- The contents of nibble 31 (sign code) must be a value in the range \(0 \times A\) to \(0 \times F\).
- The contents of each nibble 0-30 must be a value in the range \(0 \times 0\) to \(0 \times 9\).

Packed decimal values with sign codes of \(0 \times A, O \times C\), \(0 \times E\), or \(0 \times F\) are interpreted as positive values.

Packed decimal values with sign codes of \(0 \times B\) or \(0 \times D\) are interpreted as negative values.

Values greater in magnitude than \(10^{16}-1\) are too large to be represented in zoned decimal format.

For \(P S=0\), do the following.
The leftmost nibble of each digit 0-14 of the zoned decimal result is set to \(0 \times 3\).

Positive zoned decimal results are returned with a sign code of \(0 \times 3\).

Negative zoned decimal results are returned with a sign code of \(0 \times 7\).

For \(P S=1\), do the following.
The leftmost nibble of each digit 0-14 of the zoned decimal result is set to \(0 \times F\).

Positive zoned decimal results are returned with a sign code of \(0 \times C\).

Negative zoned decimal results are returned with a sign code of \(0 \times D\).

For each integer value i from 0 to 15 , do the following. The rightmost nibble of each digit \(i\) of the zoned decimal result is set to the contents of nibble \(i+15\) of src .

The result is placed into VR[ VRT].
CR field 6 is set to reflect src compared to zero, including whether or not src is too large to be represented in zoned decimal format.

If src is an invalid encoding of a packed decimal value, the contents of VR[VRT] are undefined and CR field 6 is set to 0b0001.

\section*{Special Registers Altered:}

CR field 6

\section*{Decimal Convert From Signed Quadword VX-form}
bcdcfsq.
VRT,VRB,PS
\begin{tabular}{|l|l|l|l|l|l|l|ll|}
\hline 4 & VRT & \multicolumn{2}{|c|}{2} & VRB & 1 & PS & & 385 \\
21 & 22 & 23 & & 31 \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailable()
```

Ox_flagg
|t_flag\&(EXTS[VR[VRB])< 0)
gt_flag \&(EXTS(VR[VVB]) > 0)
eq_flag\&(EXTS[VR[VRB]) = 0)
if ox_flag=0 then
result \& ConvertSI IoBCD( EXTS(VR[ VRB]) , PS)
else
result \& OxUUUU_UUUU_UUUU_UUUZ_UUUZ_UUUZ_UUUU_UUUU
VR[VRT] \&ox_flag? undefined: result
CR.bit[56] \& It_flag
CR.bit[57] \&gt_flag
CR.bit[58] \& eq_flag
CR.bit[59]}\leftarrow0\mp@subsup{X}{_}{-flag

```

Let src be the signed integer value in VR[VRB].
\(\operatorname{src}\) is placed into VR[VRT] in signed packed decimal format.

For \(P S=0\), the contents of nibble element 31 (i.e., sign code) of VR[VRT] are set to \(0 \times C\) for values greater than or equal to 0 and to \(0 \times D\) for values less than 0 .

For \(P S=1\), the contents of nibble element 31 (i.e., sign code) of VR[VRT] are set to \(0 \times F\) for values greater than or equal to 0 and to \(0 \times D\) for values less than 0 .

If the signed integer value in VR[VRB] is greater than \(10^{31} \cdot 1\) or less than \(10^{31}\). 1 , the value is too large to be represented in packed decimal format, and the contents of VR[ VRT] are undefined.

CR field 6 is set to reflect src compared to zero and whether or not \(\operatorname{sic}\) is too large in magnitude to be represented in packed decimal format.

\section*{Special Registers Altered:}

CR field 6

\section*{Decimal Convert To Signed Quadword VX-form}
bcdctsq. VRT,VRB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 04 & \({ }_{6}\) VRT & 11 & 0 & \({ }_{16}\) VRB &  & 385 \\
\hline
\end{tabular}
if MSR. VEC=O then Vector_Unavailablel)
```

inv_flag\& (VR[VRb], nibble[31] < OxA)
do i =0 to 30
inv_flag\&inv_flag| (VR[VRB],nibble[i] > Ox9)
end
src_sign \& (VR[VRB].nibble[31] = 0xB) |
(VR[VVB], ni bble[31] = OxD)
eq_flag \& (VR[VRB].nibble[0: 30] = 0)
It_flag \leftarrow(eq_flag=0) \& (src_sign=1)
gt_flag \leftarrow(eq_flag=0)\&(src_sign=0)
result \& Chop(ConvertBCDtoSI(VR[VRB]), 128)
VR[VRT] \&inv_flag? undefined: result
CR.bit[56] \& inv_flag ? Obo:It_flag
CR.bit[57] \& inv_flag? ObO: gt_flag
CR.bit[58]}\leftarrow\mathrm{ inv_flag? ObO: eq_flag
CR.bit[59] \& inv_flag

```

Let \(s i c\) be the packed decimal value in VR[VRB].
src is placed into VR[ VRT] in signed integer format.
A valid encoding of a signed packed decimal value requires the following.
- The contents of nibble 31 (sign code) must be a value in the range \(0 \times A\) to \(0 \times F\).
- The contents of each nibble 0-30 must be a value in the range \(0 \times 0\) to \(0 \times 9\).

Packed decimal values with sign codes of \(0 \times A, O \times C\), \(0 \times E\), or \(0 \times F\) are interpreted as positive values.

Packed decimal values with sign codes of \(0 \times B\) or \(0 \times D\) are interpreted as negative values.

CR field 6 is set to reflect src compared to zero.
If src is an invalid encoding of a packed decimal value, the contents of VR[ VRT] are undefined and CR field 6 is set to 0b0001.

\section*{Special Registers Altered:}

CR field 6

Vector Multiply-by-10 Unsigned Quadword VX-form
vmul10uq VRT,VRA
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{1}{|c|}{ VRT } & \multicolumn{1}{|c|}{ VRA } & \multicolumn{2}{|c|}{\(/ / /\)} & 513 & \\
0 & & 11 & 16 & & 21 & \\
\hline
\end{tabular}
if MSR.VEC=O then Vector Unavailable()
\(\operatorname{src} \leftarrow\) EXTZ \((V R[V R A])\)
\(\operatorname{prod} \leftarrow(\operatorname{src} \ll 3)+(\operatorname{src} \ll 1)\)
VR[VRT] \(\leftarrow\) Chop(prod, 128)
Let \(\operatorname{src}\) be the unsigned integer value in VR[VRA].
The rightmost 128 bits of the product of src multiplied by the value 10 are placed into VR[ VRT].

\section*{Special Registers Altered:}

None

\section*{Vector Multiply-by-10 \& write Carry Unsigned Quadword VX-form}
vmul10cuq VRT,VRA
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & \multicolumn{2}{|c|}{ VRT } & VRA & \multicolumn{2}{|c|}{ III } & \\
\hline 0 & & & 11 & 1 & 31 \\
\hline
\end{tabular}

> if MSR. VEC=0 then Vector_Unavailable|l
> sic \& EXXZ(VRR VRA])
> prod \& (sic <<3) +(sic <<1) VR[ VRT] \& Chop(prod>>128, 128)

Let src be the unsigned integer value in VR[VRA].

The product of \(\operatorname{src}\) multiplied by the value 10 is shifted right by 128 bits. The rightmost 128 bits of the shifted result is placed into VR[ VRT] .

\section*{Special Registers Altered:}

None

\section*{Vector Multiply-by-10 Extended Unsigned Quadword VX-form}
vmul10euq VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & & 577 & \\
\hline 0 & & 6 & & 11 & & \\
\hline
\end{tabular}
if MSR. VEC=0 then Vector_Unavailable el)
sic \& EXTZ[VR[VRA])
cin \(\leftarrow\) EXTz (VR[VRB], bit [124: 127])
prod \& (sic <<3) \(+(\) sic \(\ll 1)+\operatorname{cin}\)
VR[ VRT] \& Chop(prod, 128)
Let src be the unsigned integer value in VR[VRA].
Let cin be the unsigned packed decimal value in bits \(124: 127\) of VR[VRB]. Values of cin greater than 9 are undefined.

The rightmost 128 bits of the sum of cin and the product of src multiplied by the value 10 are placed into VR[ VRT].

\section*{Special Registers Altered: \\ None}

\section*{Vector Multiply-by-10 Extended \& write} Carry Unsigned Quadword VX-form
vmul10ecuq VRT,VRA,VRB
\begin{tabular}{|l|l|l|l|l|lll|}
\hline 4 & VRT & VRA & VRB & & 65 & 31 \\
\hline
\end{tabular}
if MSR.VEC=0 then Vector_Unavailable()
STC \& EXTZ(VR[VRA])
cin \(\leftarrow\) EXTZ (VRR VRB], bit [124: 127]]
prod \(\leftarrow(\operatorname{sic} \ll 3)+(\) sic \(\ll 1)+\operatorname{cin}\)
VR[ VRT] \(\leftarrow\) Chop (prod>>128, 128)
Let src be the unsigned integer value in VR[VRA].
Let cin be the unsigned packed decimal value in bits 124:127 of VR[VRA]. Values of cin greater than 9 are undefined.

The sum of cin and the product of src multiplied by the value 10 is shifted right by 128 bits. The rightmost 128 bits of the shifted result is placed into VR[ VRT] .

\section*{Special Registers Altered: \\ None}

\subsection*{6.17.3 Decimal Integer Sign Manipulation Instructions}
Decimal Copy Sign VX-form
bcdcpsgn.
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 4 & VRT,VRA,VRB \\
\hline 0 & VRT & VRA & VRB & & 833 & 31 \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailable()
inv_flag \(\leftarrow(\) VR[VRA]. nibble[31]<OxA) \(\mid\)
(VR[VRB]. ni bble[31] < OxA)
do \(\mathrm{i}=0\) to 30
inv_flag \(\leftarrow\) inv_flag \(\mid\) (VR[VRA]. nibble[i] \(>\) Oxg)
| (VR[VRB].nibble[i] >OX9)
end
src_sign \(\leftarrow(\) VR[VRB].nibble[31] \(=\) OXB) \(\mid\)
(VR[VRB].nibble[31] \(=0 \times D\) )
eq_flag \(\leftarrow(\) VR[VRA].nibble[0:30] \(=0)\)
It_flag \(\leftarrow\left(e q_{-} f a g=0\right) \&\left(s i c \_\right.\)sign=1)
gt_flag \(\leftarrow\left(e q_{\_} f a g=0\right) \&\left(\operatorname{sic}_{1} \operatorname{sign}=0\right)\)
result. nibble[0:30] \(\leftarrow\) VR[VRA], nibble[0:30]
result. nibble[31] \(\leftarrow\) VR[VRB], nibble[31]
VR[VRT] \(\leftarrow\) inv_flag ? undefined: result
CR. bit[56] \(\leftarrow\) inv_flag? ObO: It_flag
CR.bit[57] \(\leftarrow\) inv_flag? ObO: gt_flag
CR.bit[58] \(\leftarrow\) inv_flag? ObO: eq_flag
\(\underline{\text { CR. bit[59] } \leftarrow \text { inv_flag }}\)
The decimal value in VR[VRA] is placed into VR[VRT] with the sign code of the decimal value in VR[VRB].

CR field 6 is set to reflect the result compared to zero.
If either the decimal value in VR[VRA] or the decimal value in VR[ VRB] is an invalid encoding, the contents of VR[VRT] are undefined and CR field 6 is set to \(0 b 0001\).

\section*{Special Registers Altered:}

CR field 6

Decimal Set Sign VX-form
bcdsetsgn. VRT,VRB,PS

if MSR.VEC=O then Vector_Unavailablell
inv_flag \(\leftarrow(\) VR[VRB], nibble[31] < OxA)
do \(\bar{i}=0\) to 30
inv_flag ャinv_flag | (VR[VRB].nibble[i] >Oxg)
end
src_sign \(\leftarrow(V R[V R B]\), nibble[31] \(=0 \times B)\)
(VR[VRB]. nibble[31] \(=0 \times D\) )
eq_flag \(\leftarrow(V R[V R B]\).nibble[0:30] \(=0)\)
It_flag \(\leftarrow\left(e q_{-} f a g=0\right) \&\left(s i c_{\_}\right.\)sign=1)
gt_flag \(\leftarrow\left(e q_{-}^{-} f(a g=0) \&\left(s c_{1}^{-} \operatorname{sign}=0\right)\right.\)
result. nibble[0:30] \(\leftarrow\) VR[VRB], nibble[0:30]
result. nibble[31] \(\leftarrow(\) src_sign=0) ? ( \((P S=0)\) ? OXC: OXF): OXD
VR[VRT] \(\leftarrow\) inv_flag? undefined: result
CR.bit[56] \(\leftarrow\) inv_flag? Obo: It_flag
CR.bit[57] \(\leftarrow\) inv_flag? ObO: gt_flag
CR.bit[58] \(\leftarrow\) inv_flag? ObO: eq_flag
CR.bit[59] \(\leftarrow\) inv_ilag
Let \(s r^{c}\) be the packed decimal value in VR[VRB].
Packed decimal values with sign codes of \(0 \times A, 0 \times C\), \(0 \times E\), or \(0 \times F\) are interpreted as positive values.

Packed decimal values with sign codes of \(0 \times B\) or \(0 \times D\) are interpreted as negative values.

If \(\operatorname{src}\) is negative, \(\operatorname{src}\) is placed into VR[VRT] with the sign code set to \(0 \times D\).

If \(S R C\) is positive and \(P S=0, S r C\) is placed into VR[VRT] with the sign code set to \(0 \times C\).

If SrC is positive and \(P S=1, \mathrm{SrC}\) is placed into VR[VRT] with the sign code set to \(0 \times F\).

CR field 6 is set to reflect src compared to zero.
If src is an invalid encoding of a packed decimal value, the contents of VR[ VRT] are undefined and CR field 6 is set to 0b0001.

\section*{Special Registers Altered:}

CR field 6

\subsection*{6.17.4 Decimal Integer Shift and Round Instructions}
```

Decimal Shift VX-form
bcds. VRT,VRA,VRB,PS

|  | 4 | VRT |  | VRA | VRB | 1 | PS |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 193 |  |  |  |

```
```

if MSR.VEC=O then Vector_Unavailable(l

```
if MSR.VEC=O then Vector_Unavailable(l
n &EXTS(VR[VRA], byte[7])
n &EXTS(VR[VRA], byte[7])
inv_flag & (VR[VRB], nibble[31] < OXA)
inv_flag & (VR[VRB], nibble[31] < OXA)
do i =0 to 30
do i =0 to 30
    inv_flag < inv_flag | (VR[VRB],nibble[i] > Oxg)
    inv_flag < inv_flag | (VR[VRB],nibble[i] > Oxg)
end
end
src_sign \leftarrow(VR[VRB],nibble[31]=OxB)
src_sign \leftarrow(VR[VRB],nibble[31]=OxB)
    (VR[VRB], nibble[31] = OxD)
    (VR[VRB], nibble[31] = OxD)
eq_flag \leftarrow(VR[VRB],nibble[0:30]=0)
eq_flag \leftarrow(VR[VRB],nibble[0:30]=0)
It_flag \leftarrow(eq_flag=0) & (sic_sign=1)
It_flag \leftarrow(eq_flag=0) & (sic_sign=1)
gt_flag}\leftarrow(e\mp@subsup{q}{_}{\prime}f(ag=0)&(src_sign=0
gt_flag}\leftarrow(e\mp@subsup{q}{_}{\prime}f(ag=0)&(src_sign=0
if (n >si O) then do || shift left
if (n >si O) then do || shift left
    shent }\leftarrow(n<32)?n:3
    shent }\leftarrow(n<32)?n:3
    src.nibble[0:30]}\leftarrowVR[VRB].nibble[0:30
    src.nibble[0:30]}\leftarrowVR[VRB].nibble[0:30
    src.nibble[31:61]}\leftarrow\mathrm{ DUP(0bOOOO,31)
    src.nibble[31:61]}\leftarrow\mathrm{ DUP(0bOOOO,31)
    result.nibble[0:30] \leftarrow src.data, ni bble[shcnt:shcnt +30]
    result.nibble[0:30] \leftarrow src.data, ni bble[shcnt:shcnt +30]
    Ox_flag < (shent > 0) & (src.nibble[0:shcnt.1] !=0)
    Ox_flag < (shent > 0) & (src.nibble[0:shcnt.1] !=0)
end
end
else do || shift right
else do || shift right
    shent }\leftarrow((7n+1)<32)?(7n+1):3
    shent }\leftarrow((7n+1)<32)?(7n+1):3
    src.nibble[0:30]}\leftarrow\mathrm{ DUP(ObOOOO,31)
    src.nibble[0:30]}\leftarrow\mathrm{ DUP(ObOOOO,31)
    src.nibble[31:61] &VR[VRB].nibble[0:30]
    src.nibble[31:61] &VR[VRB].nibble[0:30]
    result.nibble[0:30] \leftarrow src.nibble[31-shcnt:61-shent]
    result.nibble[0:30] \leftarrow src.nibble[31-shcnt:61-shent]
    Ox_flag \leftarrow ObO
    Ox_flag \leftarrow ObO
end
end
result.nibble[31]\leftarrow(src_sign=0) ? ((PS=0) ? OxC: OxF): OxD
VR[VRT] \leftarrowinv_flag? undefined: result
CR.bit[56]\leftarrowinv_flag? ObO:It_flag
CR.bit[57] &inv_flag? ObO: gt flag
CR.bit[58] \leftarrowinv_flag? ObO: eq_flag
CR.bit[59]}\leftarrow inv_flag| Ox_flag
```

Let n be the signed integer value in byte element 7 of VR[VRA].

Let $s r^{c}$ be the signed packed decimal value in VR[VRB].
A valid encoding of a signed packed decimal value requires the following.

- The contents of nibble 31 (sign code) must be a value in the range $0 \times A$ to $0 \times F$.
- The contents of each nibble 0-30 must be a value in the range $0 \times 0$ to $0 \times 9$.

Packed decimal source operands with sign codes of $0 \times A, 0 \times C, 0 \times E$, or $0 \times F$ are interpreted as positive values.

Packed decimal source operands with sign codes of $0 \times B$ or $0 \times D$ are interpreted as negative values.

If $n$ is greater than zero, $\operatorname{src}$ is shifted left $n$ digits. Zeros are supplied to vacated digits on the right. If any non-zero digits are shifted out, an overflow occurs.

If $n$ is less than zero, src is shifted right $\cdot \mathrm{n}$ digits. Zeros are supplied to vacated digits on the left.

If the packed decimal value in VR[VRB] is negative, the sign code of the result is set to $0 b 1101$.

If the packed decimal value in VR[VRB] is positive, the sign code of the result is set to $0 b 1100$ if $P S=0$ and is set to 0 b1111 if $\mathrm{PS}=1$.

The shifted result is placed into VR[ VRT].
CR field 6 is set to reflect src compared to zero, including whether or not significant digits were shifted out when the shift count is positive (i.e., left shift operation).

If $s i c$ is an invalid encoding of a packed decimal value, the contents of VR[ VRT] are undefined and CR field 6 is set to 0 b0001.

Special Registers Altered:
CR field 6

## Decimal Unsigned Shift VX-form

| bcdus. VRT,VRA,VRB |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| - 4 | ${ }_{6} \text { VRT }$ | ${ }_{11}$ VRA | VRB <br> 16 | 1 1  <br> 21 22 23 | 129 |

if MSR.VEC=0 then Vector_Unavailablell
$n \leftarrow$ EXTS (VR[VRA] , byte[7])
$i n v \_f l a g \leftarrow 0$
do $i=0$ to 31
inv_flag $\leftarrow$ inv_flag $\mid$ (VR[VRB], nibble[i] >Oxg)
end
eq_flag $\leftarrow(V R[V R B]$.nibble[0:31] $=0)$
gt_flag $\leftarrow\left(e q_{-} f l a g=0\right)$
if $\left(\mathrm{n}>_{\text {si }} 0\right)$ then do $|\mid$ shift left
shent $\leftarrow(n<33) ? n: 32$
src.nibble[0:31] $\leftarrow$ VR[VRB]
src.nibble[32:63] $\leftarrow \operatorname{DUP}(0 b 0000,32)$
result $\leftarrow$ src.nibble[shent:shcnt +31$]$
ox_flag $\leftarrow($ shcnt $>0) \&($ src. nibble[ $0:$ shcnt -1$]!=0)$
end
else do || shift right
shent $\leftarrow((\neg n+1)<33) ?(\neg n+1): 32$
src.nibble[0:31] $\leftarrow$ DUP(ObOOOO,32)
src.nible e[32:63] $\leftarrow$ VR[VRB]
result $\leftarrow$ src. nibble[32-shent: 63 -shcnt]
$0 x+f a g \leftarrow 0$
end
VR[VRT] $\leftarrow$ inv_flag ? undefined: result
CR. bit[56] $\leftarrow$ ObO
CR.bit[57] $\leftarrow$ inv_flag? ObO: gt_flag
CR.bit[58] $\leftarrow$ inv_flag? ObO: eq_flag
CR.bit[59] $\leftarrow$ inv_-flag $\mid$ ox_flag

Let $n$ be the signed integer value in byte element 7 of VR[VRA].

Let $\operatorname{src}$ be the unsigned packed decimal value in VR[VRB].

A valid encoding of an unsigned packed decimal value requires the contents of each nibble 0-31 must be a value in the range $0 \times 0$ to $0 \times 9$.

If $n$ is greater than zero, $s i c$ is shifted left $n$ digits. Zeros are supplied to vacated digits on the right. If any non-zero digits are shifted out, an overflow occurs.

If $n$ is less than zero, $s r c$ is shifted right $\cdot n$ digits. Zeros are supplied to vacated digits on the left.

The shifted result is placed into VR[ VRT].
CR field 6 is set to reflect src compared to zero, including whether or not significant digits were shifted out when the shift count is positive (i.e., left shift operation).

If src is an invalid encoding of a packed decimal value, the contents of VR[ VRT] are undefined and CR field 6 is set to 0b0001.

Special Registers Altered:
CR field 6

## Decimal Shift and Round VX-form

bcdsr. VRT,VRA,VRB,PS

| 04 | ${ }_{6} \text { VRT }$ | ${ }_{11}$ VRA | ${ }_{16}$ VRB | $\left\|\begin{array}{c\|c}1 \\ 21\end{array}\right\| \begin{aligned} & \text { PS } \\ & 22\end{aligned}{ }_{23}$ | 449 |
| :---: | :---: | :---: | :---: | :---: | :---: |

```
if MSR.VEC=O then Vector_Unavailable()
n}\leftarrow\mathrm{ EXTS(VR[VRA], byte[7])
inv_flag \leftarrow(VR[VRB], nibble[31] < OXA)
do i =0 to 30
    inv_flag < inv_flag | (VR[VRB],nibble[i] > Oxg)
end
src_sign & (VR[VRB], nibble[31]=OxB)
    [VR[VRB], nibble[31] = OxD\
eq_flag \leftarrow(VR[VRB],nibble[0; 30]=0)
It_flag \leftarrow(eq_flag=0)&(src_sign=1)
gt_flag}\leftarrow(e\mp@subsup{q}{_}{\prime}f(ag=0)&(src_sign=0
if (n >si O) then do || shift left
    shcnt \leftarrowClamp(n, 0, 31)
    src.nibble[0:30]}\leftarrow\mathrm{ VR[VRB].nibble[0:30]
    src.nibble[31:61]}\leftarrow\mathrm{ DUP(ObOOOO,31)
    result.nibble[0:30] & src.nibble[shent:shent +30]
    ox_flag \leftarrow(shcnt > 0) & (src.nibble[0:shcnt.1] !=0)
    g_ flag}\leftarrow
end
else do |l shift right
    shont \leftarrowClamplan + 1, O, 31)
    src.nibble[0:30]}\leftarrow\mathrm{ DUP(ObOOOO,31)
    src.nibble[31:61] &VR[VRB].nibble[0:30]
    result.nibble[0:30]}\leftarrow\mathrm{ src.nibble[31-shcnt:61-shont]
    ox_flag}\leftarrow
    g_flag}\leftarrow(\mathrm{ shent > 0) & (src.nibble[62.shcnt] >ui 5)
end
result.nibble[31] \leftarrow(src_sign=0) ? ((PS=0) ? OxC: OxF): OxD
result &(g_flag=0) ? result : result tocd l
VR[VRT] &inv_flag ? undefined: result
CR.bit[56]\leftarrowinv_flag? ObO:It flag
CR.bit[57] \leftarrow inv_flag ? ObO: gt_flag
CR.bit[58]\leftarrowinv_flag? ObO: eq_flag
CR.bit[59]}\leftarrow inv_flag| Ox_flag
```

Let $n$ be the signed integer value in byte element 7 of VR[ VRA].

Let $s r c$ be the signed packed decimal value in VR[ VRB].
A valid encoding of a signed packed decimal source operand requires the following.

- The contents of nibble 31 (sign code) must be a value in the range $0 \times A$ to $0 \times F$.
- The contents of each nibble 0-30 must be a value in the range $0 \times 0$ to $0 \times 9$.

Packed decimal source operands with sign codes of $0 \times A, O \times C, O \times E$, or $0 \times F$ are interpreted as positive values.

Packed decimal source operands with sign codes of $0 \times B$ or $0 \times D$ are interpreted as negative values.

If $n$ is greater than zero, $\operatorname{src}$ is shifted left $n$ digits. Zeros are supplied to vacated digits on the right. If any non-zero digits are shifted out, an overflow occurs.

If $n$ is less than zero, src is shifted right $\cdot \mathrm{n}$ digits. Zeros are supplied to vacated digits on the left. If the value of the last digit shifted out on the right was greater than 5 , the result is incremented by 1 .

If $s r c$ is negative, the sign code of the result is set to Ob1101.

If $\operatorname{src}$ is positive, the sign code of the result is set to $0 b 1100$ if $\mathrm{PS}=0$ and is set to $0 b 1111$ if $\mathrm{PS}=1$.

The shifted and rounded result is placed into VR[ VRT].
CR field 6 is set to reflect src compared to zero, including whether or not significant digits were shifted out when the shift count is positive (i.e., left shift operation).

If src is an invalid encoding of a packed decimal value, the contents of VR[VRT] are undefined and CR field 6 is set to 0 b0001.

## Special Registers Altered:

CR field 6

### 6.17.5 Decimal Integer Truncate Instructions

Decimal Truncate VX-form
bcdtrunc. VRT,VRA,VRB,PS

| 04 | $\sigma_{6} \text { VRT }$ | ${ }_{11} \text { VRA }$ | ${ }_{16} \text { VRB }$ | $\begin{array}{\|c\|c\|c\|c\|c\|c\|c\|c\|} \hline 21 & \text { PS } & 22 \end{array}{ }_{23}$ | 257 |
| :---: | :---: | :---: | :---: | :---: | :---: |

if MSR.VEC=O then Vector_Unavailablel)
inv_flag $\leftarrow$ (VR[VRB].nibble[31]<OXA)
do $i=0$ to 30
inv_flag $\leftarrow$ inv_flag $\mid$ (VR[VRB]. nibble[i] $>0 \times 9$ )
end
length $\leftarrow$ VR[VRA], bit[48:63]
$0 x \_f a g \leftarrow 0$
src_sign $\leftarrow($ VR[VRB].nibble[31] $=0 \times B) \mid$
(VR[VRB]. nibble[31] $=0 \times D$ )
eq_flag $\leftarrow(V R[V R B]$.nibble[0:30] $=0)$
It_flag $\leftarrow$ src_sign \& req_flag
gt_flag $\leftarrow$-src_sign $\&$ req_flag
if length < 31 then do
do $\mathrm{i}=0$ to $30 \cdot$ - ength
if VR[VRB], nibble [i]! $=0 b 0000$ then ox_flag $\leftarrow 1$
result.nibble[i] $<0 b 0000$
end
if length $>0$ then do
do $i=31 \cdot \mid$ ength to 30
result.nibble[i] $\leftarrow$ VR[VRB].nibble[i] end
end
end
else result. nibble[0:30] $\leftarrow$ VR[VRB], nibble[ $0: 30]$
result. nibble[31] $\leftarrow($ src_sign=0) ? $((P S=0) ? O \times C: O X F): O X D$
VR[VRT] $\leftarrow$ inv_flag ? undefined: result
CR.bit[56] $\leftarrow$ inv_flag? ObO: It_flag
CR. bit [57] $\leftarrow$ inv_flag? ObO: gt_flag
CR.bit[58] $\leftarrow$ inv_flag? ObO: eq_flag
CR.bit[59] $\leftarrow$ inv_flag $\mid$ ox_flag

Let length be the integer value in bits 48:63 of VR[VRA].
Let $s r c$ be the signed decimal value in VR[VRB].
A valid encoding of a packed decimal source operand requires the following.

- The contents of nibble 31 (sign code) must be a value in the range $0 \times A$ to $0 \times F$.
- The contents of each nibble 0-30 must be a value in the range $0 \times 0$ to $0 \times 9$.

Packed decimal values with sign codes of $0 \times A, O \times C$, $0 \times E$, or $0 \times F$ are interpreted as positive values.

Packed decimal values with sign codes of $0 \times B$ or $0 \times D$ are interpreted as negative values.

If $s r c$ is negative, the sign code of the result is set to ob1101.

If SrC is positive, the sign code of the result is set to $0 b 1100$ if $P S=0$ and is set to $0 b 1111$ if $P S=1$.
$\operatorname{src}$ is copied into VR[VRT] with the leftmost 31-|ength digits each set to $0 b 0000$. If any of the leftmost 31-I ength digits of the signed decimal value in VR[VRB] are non-zero, an overflow occurs.

CR field 6 is set to reflect $\operatorname{src}$ compared to zero, including whether or not significant digits were truncated.

If $\operatorname{src}$ is an invalid encoding of a packed decimal value, the contents of VR[VRT] are undefined and $C R$ field 6 is set to 0b0001.

## Special Registers Altered:

CR field 6

## Decimal Unsigned Truncate VX-form

bcdutrunc. VRT,VRA,VRB

| 4 | VRT |  | VRA | VRB | 1 | 1 |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 21 | 321 |  |  |  |  |  |  |
| 0 |  |  |  | 11 | 22 | 23 |  |

```
if MSR.VEC=O then Vector_Unavailable()
inv flag<o
do i = 0 to 31
    inv_flag < inv_flag | (VR[VRB],nibble[i] > Oxg)
end
length}\leftarrowVR[VRA], bit[48:63
Ox_flag\leftarrow0
eq_flag \leftarrow(VR[VRB], nibble[0;31] = 0)
gt_flag \leftarrow(VR[VRB], nibble[0; 31]!=0)
if length < 32 then do
    do i = 0 to 31.length
        if VR[VRB],nibble[i]!=0b0000 then ox_flag \leftarrow1
        result.nibble[i]}\leftarrow0booo
    end
    if length > O then do
        do i = 32.|ength to 31
                    result.nibble[i] & VR[VRB].nibble[i]
            end
    end
end
else result }\leftarrowVR[VRB
VR[VRT] \leftarrowinv_flag?undefined: result
CR.bit[56] \leftarrowObO
CR.bit[57] \leftarrowinv_flag? ObO: gt_flag
CR.bit[58]\leftarrowinv_flag? ObO: eq_flag
CR.bit[59] \leftarrowinv_flag| Ox_flag
```

Let | engt $h$ be the integer value in bits 48:63 of VR[ VRA].
Let $s r^{c}$ be the unsigned decimal value in VR[VRB].
A valid encoding of a packed decimal source operand requires the contents of each nibble 0-31 must be a value in the range $0 \times 0$ to $0 \times 9$.
$\operatorname{src}$ is copied into VR[VRT] with the leftmost 32 . I ength digits each set to $0 b 0000$. If any of the leftmost $32 \cdot \mid$ ength digits of the signed decimal value in VR[VRB] are non-zero, an overflow occurs.

CR field 6 is set to reflect src compared to zero, including whether or not significant digits were truncated.

If $S r C$ is an invalid encoding of a packed decimal value, the contents of VR[VRT] are undefined and CR field 6 is set to 0b0001.

## Special Registers Altered:

CR field 6

### 6.18 Vector Status and Control Register Instructions

## Move To Vector Status and Control

 Register VX-form
## mtvscr VRB

| 4 | I/I |  |  | III | ${ }_{16}$ VRB |  | 1604 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  |  |  |  |  |

$$
\mathrm{VSCR} \leftarrow(\mathrm{VRB})_{96: 127}
$$

The contents of word element 3 of VRB are placed into the VSCR.

Special Registers Altered:
None

Move From Vector Status and Control Register VX-form

## mfvscr VRT

| 4 | VRT |  | I/I |  | I/I |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  |  | 15 | 1540 |  |  |  |

[^0]The contents of the VSCR are placed into word element 3 of VRT.

The remaining word elements in VRT are set to 0 .
Special Registers Altered:
None

## Chapter 7. Vector-Scalar Floating-Point Operations

### 7.1 Introduction

### 7.1.1 Overview of the Vector-Scalar Extension

Vector-Scalar Extension (VSX) provides facilities supporting vector and scalar binary floating-point operations. The following VSX features are provided to increase opportunities for vectorization.

- A unified register file, a set of Vector-Scalar Registers (VSR), supporting both scalar and vector operations is provided, eliminating the overhead of vector-scalar data transfer through storage.
- Support for word-aligned storage accesses for both scalar and vector operations is provided.
- Robust support for IEEE-754 for both vector and scalar floating-point operations is provided.
Combining the Floating-Point Registers (FPR) defined in Chapter 4. Floating-Point Facility and the Vector Registers (VR) defined in Chapter 6. Vector Facility provides additional registers to support more aggressive compiler optimizations for both vector and scalar operations.


### 7.1.1.1 Compatibility with Float-ing-Point and Decimal Floating-Point Operations

The instruction sets defined in Chapter 4. Floating-Point Facility and Chapter 5. Decimal Floating-Point retain their definition with one primary difference. The FPRs are mapped to doubleword element 0 of VSRs 0-31. The contents of doubleword 1 of the VSR corresponding to a source FPR specified by an instruction are ignored. The contents of doubleword 1 of a VSR corresponding to the target FPR specified by an instruction are undefined.

## Programming Note

Application binary interfaces extended to support VSX require special care of vector data written to VSRs 0-31 (i.e., VSRs corresponding to FPRs). Legacy scalar function calls employ doubleword-based loads and stores to preserve the contents of any nonvolatile registers, This has the adverse effect of not preserving the contents of doubleword 1 of these VSRs.

### 7.1.1.2 Compatibility with Vector Operations

The instruction set defined in Chapter 6. Vector Facility, retains its definition with one primary difference. The VRs are mapped to VSRs 32-63.

### 7.2 VSX Registers

### 7.2.1 Vector-Scalar Registers

Sixty-four 128-bit VSRs are provided. See Figure 108 All VSX floating-point computations and other data manipulation are performed on data residing in Vector-Scalar Registers, and results are placed into a VSR.

Depending on the instruction, the contents of a VSR are interpreted as a sequence of equal-length elements (words or doublewords) or as a quadword. Each of the elements is aligned within the VSR, as shown in Figure 108. Many instructions perform a
given operation in parallel on all elements in a VSR. Depending on the instruction, a word element can be interpreted as a signed integer word (SW), an unsigned integer word (UW), a logical mask value (MW), or a single-precision floating-point value (SP); a doubleword element can be interpreted as a doubleword signed integer (SD), a doubleword unsigned integer (UD), a doubleword mask (DM), or a double-precision floating-point value (DP). In the instructions descriptions, phrases like signed integer word element are used as shorthand for word element, interpreted as a signed integer.

Load and Store instructions are provided that transfer a byte, halfword, word, doubleword, or quadword between storage and a VSR.

| VSR[0] |  |
| :---: | :---: |
| VSR[1] |  |
| $\ldots$ |  |
| $\ldots$ |  |
| 0 | VSR[62] |
|  | VSR[63] |

Figure 108.Vector-Scalar Registers


| SQ/UQ/QP/BCD |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| SD/UD/MD/DP 0 |  |  |  | SD/UD/MD/DP 1 |  |  |  |
| SW/UW/MW/SP 0 |  | SW/UW/MW/SP 1 |  | SW/UW/MW/SP 2 |  | SW/UW/MW/SP 3 |  |
| HP 0 | HP 1 | HP 2 | HP 3 | HP 4 | HP 5 | HP 6 | HP 7 |
|  |  |  |  |  |  |  |  |

Figure 109.Vector-Scalar Register Elements

### 7.2.1.1 Floating-Point Registers

Chapter 4. Floating-Point Facility provides 32 64-bit FPRs. Chapter 5. Decimal Floating-Point also employs FPRs in decimal floating-point (DFP) operations. When VSX is implemented, the 32 FPRs are mapped to doubleword 0 of VSRs $0-31$. For example, $\operatorname{FPR}[0]$ is located in doubleword element 0 of $\operatorname{VSR}[0]$, $\operatorname{FPR}[1]$ is located in doubleword element 0 of VSR[1], and so forth.

All instructions that operate on an FPR are redefined to operate on doubleword element 0 of the corresponding VSR. The contents of doubleword element 1 of the VSR corresponding to a source FPR or FPR pair for these instructions are ignored and the contents of doubleword element 1 of the VSR corresponding to the target FPR or FPR pair for these instructions are undefined.


Figure 110.Floating-Point Registers as part of VSRs

## Version 3.0

### 7.2.1.2 Vector Registers

Chapter 6. Vector Facility provides 32 128-bit VRs. When VSX is implemented, the 32 VRs are mapped to VSRs 32-63. For example, VR[0] is located in $\operatorname{VSR}[32], \operatorname{VR}[1]$ is located in $\operatorname{VSR}[33]$, and so forth.

All instructions that operate on a VR are redefined to operate on the corresponding VSR.


Figure 111.Vector Registers as part of VSRs

### 7.2.2 Floating-Point Status and Control Register

The Floating-Point Status and Control Register (FPSCR) controls the handling of floating-point exceptions and records status resulting from the floating-point operations. Bits 0:19 and 32:55 are status bits. Bits 56:63 are control bits.

The exception status bits in the FPSCR (bits 35:44, $53: 55$ ) are sticky; that is, once set to 1 they remain set to 1 until they are set to 0 by an mcrfs, mtfsfi, mtfsf, or mtfsb0 instruction. The exception summary bits in the FPSCR (FX, FEX, and VX, which are bits 32:34) are not considered to be "exception status bits", and only FX is sticky.

## Programming Note

Access to Move To FPSCR and Move From $F P S C R$ instructions requires $\mathrm{FP}=1$.

FEX and VX are simply the ORs of other FPSCR bits. Therefore these two bits are not listed among the FPSCR bits affected by the various instructions.

The bit definitions for the FPSCR are as follows.

## Bits Definition

0:28 Decimal Floating-Point Rounding Control (DRN)
This field is not used by VSX instructions.
32 Floating-Point Exception Summary (FX)
Every floating-point instruction, except mtfsfi and mtfsf, implicitly sets FX to 1 if that instruction causes any of the floating-point exception bits in the FPSCR to change from 0 to 1. mcrfs, mtfsfi, mtfsf, mtfsbo, and mtfsb1 can alter FX explicitly.

[^1]
## Bits Definition

Floating-Point Invalid Operation Exception Summary (VX)
This bit is the OR of all the Invalid Operation exception bits. mcrfs, mtfsfi, mtfsf, mtfsbO, and $\boldsymbol{m t f s} b 1$ cannot alter VX explicitly.

Floating-Point Overflow Exception (OX)
This bit is set to 1 when a VSX Scalar Floating-Point Arithmetic, VSX Vector Floating-Point Arithmetic, VSX Scalar DP-SP Conversion or VSX Vector DP-SP Conversion class instruction causes an Overflow exception. See Section 7.4.3, "Floating-Point Overflow Exception" on page 406.
This bit can be set to 0 or 1 by a Move To FPSCR class instruction.

Floating-Point Underflow Exception (UX)
This bit is set to 1 when a VSX Scalar Floating-Point Arithmetic, VSX Vector Floating-Point Arithmetic, VSX Scalar DP-SP Conversion or VSX Vector DP-SP Conversion class instruction causes an Underflow exception. See Section 7.4.4, "Floating-Point Underflow Exception" on page 411.
This bit can be set to 0 or 1 by a Move To FPSCR class instruction.

Floating-Point Zero Divide Exception (ZX) This bit is set to 1 when a VSX Scalar Floating-Point Arithmetic or VSX Vector Floating-Point Arithmetic class instruction causes an Zero Divide exception. See Section 7.4.2, "Floating-Point Zero Divide Exception" on page 403.
This bit can be set to 0 or 1 by a Move To FPSCR class instruction.

Floating-Point Inexact Exception (XX)
This bit is set to 1 when a VSX Scalar Floating-Point Arithmetic, VSX Vector Floating-Point Arithmetic, VSX Scalar Integer Conversion, VSX Vector Integer Conversion, VSX Scalar Round to Floating-Point Integer, or VSX Vector Round to Floating-Point Integer class instruction causes an Inexact exception. See Section 7.4.5, "Floating-Point Inexact Exception" on page 416.
This bit can be set to 0 or 1 by a Move To FPSCR class instruction.
$41 \quad$ Floating-Point Invalid Operation Exception (InfㄴInf) (VXIDI)
This bit is set to 1 when a VSX Scalar Floating-Point Arithmetic and VSX Vector Floating-Point Arithmetic class instruction causes an Infinity $\div$ Infinity type Invalid Operation exception. See Section 7.4.1, "Floating-Point Invalid Operation Exception" on page 392.

This bit can be set to 0 or 1 by a Move To FPSCR class instruction.

42 Floating-Point Invalid Operation Exception (Zero $\div$ Zero) (VXZDZ)
This bit is set to 1 when a VSX Scalar Floating-Point Arithmetic and VSX Vector Floating-Point Arithmetic class instruction causes a Zero $\div$ Zero type Invalid Operation exception. See Section 7.4.1 , "Floating-Point Invalid Operation Exception" on page 392.

This bit can be set to 0 or 1 by a Move To FPSCR class instruction.

## Definition

Floating-Point Invalid Operation Exception (Inf×Zero) (VXIMZ)
This bit is set to 1 when a VSX Scalar Floating-Point Arithmetic and VSX Vector Floating-Point Arithmetic class instruction causes a Infinity $\times$ Zero type Invalid Operation exception. See Section 7.4.1, "Floating-Point Invalid Operation Exception" on page 392.

This bit can be set to 0 or 1 by a Move To FPSCR class instruction.

Floating-Point Invalid Operation Exception (Invalid Compare) (VXVC)
This bit is set to 1 when a VSX Scalar Compare Double-Precision, VSX Vector Compare Double-Precision, or VSX Vector Compare Single-Precision class instruction causes an Invalid Compare type Invalid Operation exception. See Section 7.4.1, "Floating-Point Invalid Operation Exception" on page 392.
This bit can be set to 0 or 1 by a Move To FPSCR class instruction.

Floating-Point Fraction Rounded (F R)
This bit is set to 0 or 1 by VSX Scalar Floating-Point Arithmetic, VSX Scalar Integer Conversion, and VSX Scalar Round to Floating-Point Integer class instructions to indicate whether or not the fraction was incremented during rounding. See Section 7.3.2.6 , "Rounding" on page 383. This bit is not sticky.

Floating-Point Fraction Inexact ( Fl )
This bit is set to 0 or 1 by VSX Scalar Floating-Point Arithmetic, VSX Scalar Integer Conversion, and VSX Scalar Round to Floating-Point Integer class instructions to indicate whether or not the rounded result is inexact or the instruction caused a disabled Overflow exception. See Section 7.3.2.6 on page 383. This bit is not sticky.

See the definition of $X X$, above, regarding the relationship between Fl and XX .

| Bits | Definition |
| :---: | :---: |
| 47:51 | Floating-Point Result Flags (FPRF) |
|  | VSX Scalar Floating-Point Arithmetic, VSX |
|  | Scalar DP-SP Conversion, VSX Scalar |
|  | Convert Integer to Double-Precision, and |
|  | VSX Scalar Round to Double-Precision |
|  | Integer class instructions set this field based |
|  | on the result placed into the target register |
|  | portion of the result is undefined then the |
|  | value placed into FPRF is undefined. |
|  | For VSX Scalar Convert Double-Precision to |
|  | Integer class instructions, the value placed into FPRF is undefined. |
|  | Additional details are as follows. |
| 47 | Floating-Point Result Class |
|  | Descriptor (C) |
|  | VSX Scalar Floating-Point Arithmetic, VSX |
|  | Scalar DP-SP Conversion, VSX Scalar |
|  | Convert Integer to Double-Precision, and |
|  | VSX Scalar Round to Double-Precision |
|  | Integer class instructions set this bit with the |
|  | FPCC bits, to indicate the class of the result as |
|  | shown in Table 2, "Floating-Point Result Flags," on page 373. |
| 48:51 | Floating-Point Condition Code (FPCC) |
|  | VSX Scalar Compare Double-Precision |
|  | instruction sets one of the FPCC bits to 1 and |
|  | the other three FPCC bits to 0 based on the relative values of the operands being |
|  | compared. |
|  | VSX Scalar Floating-Point Arithmetic, VSX |
|  | Scalar DP-SP Conversion, VSX Scalar |
|  | Convert Integer to Double-Precision, and |
|  | VSX Scalar Round to Double-Precision |
|  | Integer class instructions set the FPCC bits with |
|  | the C bit, to indicate the class of the result as |
|  | shown in Table 2, "Floating-Point Result |
|  | Flags," on page 373. Note that in this case |
|  | the high-order three bits of the FPCC retain |
|  | their relational significance indicating that the |
|  | value is less than, greater than, or equal to zero. |
| 48 | Floating-Point Less Than or Negative (FL) |
| 49 | Floating-Point Greater Than |
|  | Positive (FG) |
| 50 | Floating-Point Equal or Zero (FE) |
| 51 | Floating-Point Unordered or NaN (FU) |

## Definition

Reserved

Floating-Point Invalid Operation Exception (Software-Defined Condition) (VXSOFT)
This bit can be altered only by mcrfs, mtfsfi, mtfsf, mtfsb0, or mtfsb1. See Section 7.4.1, "Floating-Point Invalid Operation Exception" on page 392.

## Programming Note

VXSOFT can be used by software to indicate the occurrence of an arbitrary, software-defined, condition that is to be treated as an Invalid Operation exception. For example, the bit could be set by a program that computes a base 10 logarithm if the supplied input is negative.

Floating-Point Invalid Operation Exception (Invalid Square Root) (VXSQRT)
This bit is set to 1 when a VSX Scalar Floating-Point Arithmetic or VSX Vector Floating-Point Arithmetic class instruction causes a Invalid Square Root type Invalid Operation exception. See Section 7.4.1, "Floating-Point Invalid Operation Exception" on page 392.
This bit can be set to 0 or 1 by a Move To FPSCR class instruction.

Floating-Point Invalid Operation Exception (Invalid Integer Convert) (VXCVI)
This bit is set to 1 when a VSX Scalar Convert Double-Precision to Integer, VSX Vector Convert Double-Precision to Integer, or VSX Vector Convert Single-Precision to Integer class instruction causes a Invalid Integer Convert type Invalid Operation exception. See Section 7.4.1, "Floating-Point Invalid Operation Exception" on page 392.
This bit can be set to 0 or 1 by a Move To FPSCR class instruction.

Floating-Point Invalid Operation Exception Enable (VE)
This bit is used by VSX Scalar Floating-Point and VSX Vector Floating-Point class instructions to enable trapping on Invalid Operation exceptions. See Section 7.4.1, "Floating-Point Invalid Operation Exception" on page 392.
Bits

57 Floating-Point Overflow Exception Enable (OE)
This bit is used by VSX Scalar Floating-Point and VSX Vector Floating-Point class instructions to enable trapping on Overflow exceptions. See Section 7.4.3, "Floating-Point Overflow Exception" on page 406.

58 Floating-Point Underflow Exception Enable (UE)
This bit is used by VSX Scalar Floating-Point and VSX Vector Floating-Point class instructions to enable trapping on Underflow exceptions. See Section 7.4.4, "Floating-Point Underflow Exception" on page 411.

59 Floating-Point Zero Divide Exception Enable (ZE)
This bit is used by VSX Scalar Floating-Point and VSX Vector Floating-Point class instructions to enable trapping on Zero Divide exceptions. See Section 7.4.2, "Floating-Point Zero Divide Exception" on page 403.

60 Floating-Point Inexact Exception Enable (XE)
This bit is used by VSX Scalar Floating-Point and VSX Vector Floating-Point class instructions to enable trapping on Inexact exceptions. See Section 7.4.5, "Floating-Point Inexact Exception" on page 416.

61 Floating-Point Non-IEEE Mode (NI)
Floating-point non-IEEE mode is optional. If floating-point non-IEEE mode is not implemented, this bit is treated as reserved, and the remainder of the definition of this bit does not apply.
If floating-point non-IEEE mode is implemented, this bit has the following meaning.
0 The processor is not in floating-point non-IEEE mode (i.e., all floating-point operations conform to the IEEE standard).
1 The processor is in floating-point non-IEEE mode.

Bits Definition
61 Floating-Point Non-IEEE Mode (NI) (continued)
When the processor is in floating-point non-IEEE mode, the remaining FPSCR bits is permitted to have meanings different from those given in this document, and floating-point operations need not conform to the IEEE standard. The effects of executing a given floating-point instruction with $\mathrm{Nl}=1$, and any additional requirements for using non-IEEE mode, are implementation-dependent. The results of executing a given instruction in non-IEEE mode is permitted to vary between implementations, and between different executions on the same implementation.

## - Programming Note

When the processor is in floating-point non-IEEE mode, the results of floating-point operations is permitted to be approximate, and performance for these operations might be better, more predictable, or less data-dependent than when the processor is not in non-IEEE mode. For example, in non-IEEE mode an implementation is permitted to return 0 instead of a denormalized number and return a large number instead of an infinity.

62:63
Floating-Point Rounding Control (RN)
This field is used by VSX Scalar Floating-Point and VSX Vector Floating-Point class instructions that round their result and the rounding mode is not implied by the opcode.
This bit can be explicitly set or reset by a new Move To FPSCR class instruction.

See Section 7.3.2.6 , "Rounding" on page 383.

00 Round to Nearest Even
01 Round toward Zero
10 Round toward +Infinity
11 Round toward -Infinity

| Result Flags |  |  |  | Result Value Class |  |
| :---: | :---: | :---: | :---: | :---: | :--- |
| $\mathbf{C}$ | FL | FG | FE | FU |  |
| 1 | 0 | 0 | 0 | 1 | Quiet NaN |
| 0 | 1 | 0 | 0 | 1 | - Infinity |
| 0 | 1 | 0 | 0 | 0 | - Normalized Number |
| 1 | 1 | 0 | 0 | 0 | - Denormalized Number |
| 1 | 0 | 0 | 1 | 0 | - Zero |
| 0 | 0 | 0 | 1 | 0 | + Zero |
| 1 | 0 | 1 | 0 | 0 | + Denormalized Number |
| 0 | 0 | 1 | 0 | 0 | + Normalized Number |
| 0 | 0 | 1 | 0 | 1 | + Infinity |

Table 2. Floating-Point Result Flags

### 7.3 VSX Operations

### 7.3.1 VSX Floating-Point Arithmetic Overview

This section describes the floating-point arithmetic and exception model supported by Vector-Scalar Extension. Except for extensions to support 32-bit single-precision floating-point vector operations, the models are identical to that described in Chapter 4. Floating-Point Facility.

The processor (augmented by appropriate software support, where required) implements a floating-point system compliant with the ANSI/IEEE Standard 754-1985, IEEE Standard for Binary Floating-Point Arithmetic (hereafter referred to as the IEEE standard). That standard defines certain required "operations" (addition, subtraction, and so on). Herein, the term, floating-point operation, is used to refer to one of these required operations and to additional operations defined (e.g., those performed by Multiply-Add or Reciprocal Estimate instructions). A Non-IEEE mode is also provided. This mode, which is permitted to produce results not in strict compliance with the IEEE standard, allows shorter latency.

Instructions are provided to perform arithmetic, rounding, conversion, comparison, and other operations in VSRs; to move floating-point data between storage and these registers.

These instructions are divided into two categories.

- computational instructions

The computational instructions are those that perform addition, subtraction, multiplication, division, extracting the square root, rounding, conversion, comparison, and combinations of these operations. These instructions provide the floating-point operations. There are two forms of computational instructions, scalar, which perform a single floating-point operation, and vector, which perform either two double-precision floating-point operations or four single-precision operations. Computational instructions place status information into the Floating-Point Status and Control Register. They are the instructions described in Sections 7.6.1.3 through 7.6.1.7.2.

- noncomputational instructions

The noncomputational instructions are those that perform loads and stores, move the contents of a VSR to another floating-point register possibly altering the sign, and select the value from one of two VSRs based on the value in a third VSR. The
operations performed by these instructions are not considered floating-point operations. These instructions do not alter the Floating-Point Status and Control Register. They are the instructions listed in Sections 7.6.1.1, 7.6.1.2.1, and 7.6.1.9 through 7.6.1.10.

A floating-point number consists of a signed exponent and a signed significand. The quantity expressed by this number is the product of the significand and the number 2 exponent. Encodings are provided in the data format to represent finite numeric values, $\pm$ Infinity, and values that are "Not a Number" ( NaN ). Operations involving infinities produce results obeying traditional mathematical conventions. NaNs have no mathematical interpretation. Their encoding permits a variable diagnostic information field. NaNs might be used to indicate such things as uninitialized variables and can be produced by certain invalid operations.

There is one class of exceptional events that occur during instruction execution that is unique to Vector-Scalar Extension and Floating-Point: the Floating-Point Exception. Floating-point exceptions are signaled with bits set in the FPSCR. They can cause the system floating-point enabled exception error handler to be invoked, precisely or imprecisely, if the proper control bits are set.

## Floating-Point Exceptions

The following floating-point exceptions are detected by the processor:

```
- Invalid Operation exception (VX)
    SNaN
    Infinity-Infinity
    Infinity\divInfinity
    Zero\divZero
    Infinity\timesZero
    Invalid Compare
    Software-Defined Condition
    Invalid Square Root
    Invalid Integer Convert (VXCVI)
- Zero Divide exception (ZX)
- Overflow exception
- Underflow exception
- Inexact exception
(VXSNAN)
(VXISI)
(VXIDI )
(VXZDZ)
(VXI MZ)
(VXVC)
(VXSOFT)
(VXSQRT)
(VXCVI)
- Inexact exception

Each floating-point exception, and each category of Invalid Operation exception, has an exception bit in the FPSCR. In addition, each floating-point exception has a corresponding enable bit in the FPSCR. See Section 7.2.2, "Floating-Point Status and Control Register" on page 369 for a description of these exception and enable bits, and Section 7.3.3, "VSX Floating-Point Execution Models" on page 386 for a detailed discussion of floating-point exceptions, including the effects of the enable bits.

\subsection*{7.3.2 VSX Floating-Point Data}

\subsection*{7.3.2.1 Data Format}

This architecture defines the representation of a floating-point value in three different binary fixed-length formats, 16 -bit half-precision, 32-bit single-precision format, 64-bit double-precision format, and 128-bit quad-precision format. The half-precision format is used for half-precision data in storage and registers. The single-precision format is used for single-precision data in storage and registers. The double-precision format is used for double-precision data in storage and registers. The quad-precision format is used for quad-precision floating-point data in storage and registers.

The lengths of the exponent and the fraction fields differ between these three formats. The structure of the half-precision, single-precision, double-precision, and quad-precision formats is shown below.

Values in floating-point format are composed of three fields:
\(5 \quad\) sign bit
EXP exponent+bias
FRACTI ON fraction

Representation of numeric values in the floating-point formats consists of a sign bit ( \(\varsigma\) ), a biased exponent (EXP), and the fraction portion (FRACTION) of the significand. The significand consists of a leading implied bit concatenated on the right with the FRACTI ON. This leading implied bit is 1 for normalized numbers
I and 0 for denormalized (subnormal) numbers or zero and is located in the unit bit position (that is, the first bit to the left of the binary point). Values representable \| within the three floating-point formats can be specified by the parameters listed in Table 3.
\begin{tabular}{l|l|l|lll|}
\hline I & S & EXP & & FRACTION & \\
\hline I & 01 & 6 & & 15
\end{tabular}
| Figure 112. Floating-point half-precision format
\begin{tabular}{|l|l|l|}
\hline\(S\) & EXP & FRACTION \\
\hline 0 & & 9
\end{tabular}

Figure 113. Floating-point single-precision format
\begin{tabular}{|ll|ll|}
\hline\(S\) & EXP & & FRACTION \\
\hline 01 & 12 & 63
\end{tabular}

Figure 114.Floating-point double-precision format
I
\begin{tabular}{|ll|ll|}
\hline\(S\) & EXP & & FRACTION \\
\hline 01 & 16 & 127 \\
\hline
\end{tabular}

Figure 115.Floating-point quad-precision format (binary128)

Version 3.0
\begin{tabular}{l}
\hline \\
\hline Exponent Bias \\
\hline Maximum Exponent (Emax) \\
\hline Minimum Exponent (Emin) \\
\hline Widths (bits): \\
Format \\
Sign \\
Exponent \\
Fraction \\
Significand
\end{tabular}

Table 3. IEEE floating-point fields

\subsection*{7.3.2.2 Value Representation}

This architecture defines numeric and nonnumeric I values representable within each of the three supported formats. The numeric values are approximations to the real numbers and include the normalized numbers, denormalized numbers, and zero values. The nonnumeric values representable are the infinities and the Not a Numbers (NaNs). The infinities are adjoined to the real numbers, but are not numbers themselves, and the standard rules of arithmetic do not hold when they are used in an operation. They are related to the real numbers by order alone. It is possible however to define restricted operations among numbers and infinities as defined below. The relative location on the real number line for each of the defined entities is shown in Figure 116.

Figure 116.Approximation to real numbers


The NaNs are not related to the numeric values or infinities by order or value but are encodings used to convey diagnostic information such as the representation of uninitialized variables.

The following is a description of the different floating-point values defined in the architecture:

\section*{Binary floating-point numbers}

Machine representable values used as approximations to real numbers. Three categories of numbers are supported: normalized numbers, denormalized numbers, and zero values.

\section*{Normalized numbers ( \(\pm\) NOR)}

These are values that have a biased exponent value in the range:

I \(\quad 1\) to 30 in half-precision format
1 to 254 in single-precision format
1 to 2046 in double-precision format
1 to 32766 in quad-precision format
They are values in which the implied unit bit is 1 . Normalized numbers are interpreted as follows:
\[
\text { NOR }=(\cdot 1)^{5} \times 2^{E} \times(1 . f r a c t i o n)
\]
where \(s\) is the sign, \(E\) is the unbiased exponent, and lefraction is the significand, which is composed of a leading unit bit (implied bit) and a fraction part.

\section*{Zero values ( \(\pm 0\) )}

These are values that have a biased exponent value of zero and a fraction value of zero. Zeros
can have a positive or negative sign. The sign of zero is ignored by comparison operations (that is, comparison regards +0 as equal to -0 ).

\section*{Denormalized numbers ( \(\pm\) DEN)}

These are values that have a biased exponent value of zero and a nonzero fraction value. They are nonzero numbers smaller in magnitude than the representable normalized numbers. They are values in which the implied unit bit is 0 . Denormalized numbers are interpreted as follows:
\[
\text { DEN }=(\cdot 1)^{s} \times 2^{\text {Emin }} \times(0, f r a c t i o n)
\]
where Emin is the minimum representable exponent value.
- 14 for half-precision
- 126 for single-precision
- 1022 for double-precision
. 16382 for quad-precision.

\section*{Infinities ( \(\pm\) INF)}

These are values that have the maximum biased exponent value:

> 31 in half-precision format
> 255 in single-precision format
> 2047 in double-precision format
> 32767 in quad-precision format
and a zero fraction value. They are used to approximate values greater in magnitude than the maximum normalized value.

Infinity arithmetic is defined as the limiting case of real arithmetic, with restricted operations defined among numbers and infinities. Infinities and the real numbers can be related by ordering in the affine sense:
-Infinity < every finite number < +Infinity

Arithmetic on infinities is always exact and does not signal any exception, except when an exception occurs due to the invalid operations as described in Section 7.4.1 , "Floating-Point Invalid Operation Exception" on page 392.

For comparison operations, +Infinity compares equal to +Infinity and -Infinity compares equal to -Infinity.

\section*{Not a Numbers (NaNs)}

These are values that have the maximum biased exponent value and a nonzero fraction value. The sign bit is ignored (that is, NaNs are neither positive nor negative). If the high-order bit of the fraction field is 0 , the NaN is a Signaling NaN ; otherwise it is a Quiet NaN .

Signaling NaNs are used to signal exceptions when they appear as operands of computational instructions.

Quiet NaNs are used to represent the results of certain invalid operations, such as invalid arithmetic operations on infinities or on NaNs, when Invalid Operation exception is disabled (VE=0). Quiet NaNs propagate through all floating-point operations except ordered comparison and conversion to integer. Quiet NaNs do not signal exceptions, except for ordered comparison and conversion to integer operations. Specific encodings in QNaNs can thus be preserved through a sequence of floating-point operations, and used to convey diagnostic information to help identify results from invalid operations.

Assume the following generic arithmetic templates.
```

$\mathrm{f}(\mathrm{src} 1, \operatorname{src} 3, \operatorname{src}(2)$
ex: result $=(\operatorname{src} 1 \times \operatorname{src} 3) \cdot \operatorname{src} 2$
$\mathrm{f}(\mathrm{src} 1, \operatorname{src} 2)$
ex: result $=\operatorname{src} 1 \times \operatorname{src} 2$
ex: result $=\operatorname{src} 1+\operatorname{src} 2$
f(srel)
ex: result $=\mathrm{f}(\mathrm{src}(1)$

```

When a QNaN is the result of a floating-point operation because one of the operands is a NaN or because a QNaN was generated due to a trap-disabled Invalid Operation exception, the following rule is applied to determine the NaN with the high-order fraction bit set to 1 that is to be stored as the result.
```

if srcl is a NaN
then result = Quiet(srcl)
else if srcz is a NaN (if there is a src2)
then result = Quiet(src2)
else if src3 is a NaN (if there is a src3)
then result = Quiet(src3)
else if disabled invalid operation exception
then result = generated QNaN

```
where Quiet ( x ) means x if x is a QNaN and x converted to a QNaN if \(x\) is an SNaN . Any instruction that generates a QNaN as the result of a disabled Invalid Operation exception generates the value,
\(0 \times 7 E 00\) for half-precision results,
\(0 \times 7\) FCO_ 0000 for single-precision results,
OX7FF8_0000_0000_0000 for double-precision results,

OX7FFF_ \(8000 \_0000 \_0000 \_0000 \_0000 \_0000 \_0000\) for quad-precision results.

Note that the M -form multiply-add-type instructions use the \(B\) source operand to specify srcs and the T target operand to specify src2, whereas A-form multiply-add-type instructions use the \(B\) source operand to specify \(\operatorname{src} 2\) and the \(T\) target operand to specify sic3.

A double-precision NaN is considered to be representable in single-precision format if and only if the low-order 29 bits of the double-precision NaN's fraction are zero.

\subsection*{7.3.2.3 Sign of Result}

The following rules govern the sign of the result of an arithmetic, rounding, or conversion operation, when the operation does not yield an exception. They apply even when the operands or results are zeros or infinities.
- The sign of the result of an add operation is the sign of the operand having the larger absolute value. If both operands have the same signs, the sign of the result of an add operation is the same as the sign of the operands. The sign of the result of the subtract operation \(x \cdot y\) is the same as the sign of the result of the add operation \(x+(\cdot y)\).

When the sum of two operands with opposite sign, or the difference of two operands with the same signs, is exactly zero, the sign of the result is positive in all rounding modes except Round toward -Infinity, in which mode the sign is negative.
- The sign of the result of a multiply or divide operation is the Exclusive OR of the signs of the operands.
- The sign of the result of a Square Root or Reciprocal Square Root Estimate operation is always positive, except that the square root of -0 is -0 and the reciprocal square root of -0 is - Infinity.
- The sign of the result of a Convert From Integer or Round to Floating-Point Integer operation is the sign of the operand being converted.

For the Multiply-Add instructions, the rules given above are applied first to the multiply operation and then to the add or subtract operation (one of the inputs to the add or subtract operation is the result of the multiply operation).

\subsection*{7.3.2.4 Normalization and Denormalization}

The intermediate result of an arithmetic instruction can require normalization and/or denormalization as described below. Normalization and denormalization do not affect the sign of the result.

When an arithmetic or rounding instruction produces an intermediate result which carries out of the significand, or in which the significand is nonzero but has a leading zero bit, it is not a normalized number and must be normalized before it is stored. For the carry-out case, the significand is shifted right one bit, with a one shifted into the leading significand bit, and the exponent is incremented by one. For the leading-zero case, the significand is shifted left while decrementing its exponent by one for each bit shifted, until the leading significand bit becomes one. The Guard bit and the Round bit (see Section 7.3.3.1, "VSX Execution Model for IEEE Operations" on page 386) participate in the shift with zeros shifted into the Round bit. The exponent is regarded as if its range were unlimited.

After normalization, or if normalization was not required, the intermediate result can have a nonzero significand and an exponent value that is less than the minimum value that can be represented in the format specified for the result. In this case, the intermediate result is said to be "Tiny" and the stored result is determined by the rules described in Section 7.4.4, "Floating-Point Underflow Exception" on page 411. These rules can require denormalization.

A number is denormalized by shifting its significand right while incrementing its exponent by 1 for each bit shifted, until the exponent is equal to the format's minimum value. If any significant bits are lost in this shifting process, "Loss of Accuracy" has occurred (See Section 7.4.4, "Floating-Point Underflow Exception" on page 411) and Underflow exception is signaled.

\section*{Engineering Note}

When denormalized numbers are operands of multiply, divide, and square root operations, some implementations might prenormalize the operands internally before performing the operations.

\subsection*{7.3.2.5 Data Handling and Precision}

Scalar double-precision floating-point data is represented in double-precision format in VSRs and storage.

Vector double-precision floating-point data is represented in double-precision format in VSRs and storage.

Scalar single-precision floating-point data is represented in double-precision format in VSRs and in single-precision format in storage.

Vector single-precision floating-point data is represented in single-precision format in VSRs and storage.

Double-precision operands may be used as input for double-precision scalar arithmetic operations.

Double-precision operands may be used as input for single-precision scalar arithmetic operations when trapping on overflow and underflow exceptions is disabled.

Single-precision operands may be used as input for double-precision and single-precision scalar arithmetic operations.

Double-precision operands may be used as input for double-precision vector arithmetic operations.

Single-precision operands may be used as input for single-precison vector arithmetic operations.

Instructions are also provided for manipulations which do not require double-precision or single-precision. In addition, instructions are provided to access an integer representation in GPRs.

\section*{Half-Precision Operands}

Instructions are provided to convert between half-precision and single-precision formats for vector data in VSRs and between half-precision and double-precision formats for scalar data. Note that scalar double-precision format is identical to scalar single-precision format.

An instruction is provided to explicitly convert half-precision format operands in a VSR to single-precision format. Scalar single-precision floating-point is enabled with six types of instruction.
1. VSX Scalar Convert Half-Precision format to Double-Precision format XX2-form

The half-precision floating-point value in the rightmost halfword of each doubleword element 0 of the source VSR is placed into the doubleword element 0 of the target VSR in double-precision format.
2. VSX Scalar round \& Convert Double-Precision format to Half-Precision format XX2-form

The double-precision value in doubleword element 0 of the source VSR is rounded to to half-precision, checking the exponent for half-precision range
and handling any exceptions according to respective enable bits,and places the result into the rightmost halfword of doubleword element 0 of the target VSR in half-precision format.

Source operand values greater in magnitude than \(2^{39}\) when Overflow is enabled ( \(0 E=1\) ) produce undefined results because the value cannot be scaled into the half-precision normalized range.

Source operand values smaller in magnitude than \(2^{.38}\) when Underflow is enabled ( \(U E=1\) ) produce undefined results because the value cannot be scaled into the half-precision normalized range.
3. VSX Vector Convert Half-Precision format to Single-Precision format XX2-form

The half-precision floating-point value in the rightmost halfword of each word element of the source VSR is placed into the corresponding word element of the target VSR in single-precision format.
4. VSX Vector round and Convert Single-Precision format to Half-Precision format XX2-form

The single-precision floating-point value in each word element \(i\) of the source VSR is rounded to half-precision and placed into the rightmost halfword of the corresponding word element of the target VSR in half-precision format.

\section*{Single-Precision Operands}

For single-precision scalar data, a conversion from single-precision format to double-precision format is performed when loading from storage into a VSR and a conversion from double-precision format to single-precision format is performed when storing from a VSR to storage. No floating-point exceptions are caused by these instructions.

Instructions are provided to convert between single-precision and double-precision formats for scalar and vector data in VSRs.

An instruction is provided to explicitly convert a double format operand in a VSR to single-precision. Scalar single-precision floating-point is enabled with six types of instruction.

\section*{1. Load Scalar Single-Precision}

This form of instruction accesses a floating-point operand in single-precision format in storage, converts it to double-precision format, and loads it into a VSR. No floating-point exceptions are caused by these instructions.

\section*{2. Scalar Round to Single-Precision}
xsrsp rounds a double-precision operand to single-precision, checking the exponent for single-precision range and handling any exceptions according to respective enable bits, and places that operand into a VSR in double-precision format. For results produced by single-precision arithmetic instructions, single-precision loads, and other instances of xsrsp, xsrsp does not alter the value. Values greater in magnitude than \(2^{319}\) when Overflow is enabled \((O E=1)\) produce undefined results because the value cannot be scaled back into the normalized range. Values smaller in magnitude than \(2^{-318}\) when Underflow is enabled \((U E=1)\) produce undefined results because the value cannot be scaled back into the normalized range.
3. Scalar Convert Single-Precision to Double-Precision
xscvspdp accesses a floating-point operand in single-precision format from word element 0 of the source VSR, converts it to double-precision format, and places it into doubleword element 0 of the target VSR.
4. Scalar Convert Double-Precision to Single-Precision
xscvdpsp rounds the double-precision floating-point value in doubleword element 0 of the source VSR to single-precision, and places the result into word element 0 of the target VSR in single-precision format. This function would be used to port scalar floating-point data to a format compatible for single-precision vector operations. Values greater in magnitude than \(2^{319}\) when Overflow is enabled ( \(O E=1\) ) produce undefined results because the value cannot be scaled back into the normalized range. Values smaller in magnitude than 2.318 when Underflow is enabled \((U E=1)\) produce undefined results because the value cannot be scaled back into the normalized range.
5. VSX Scalar Single-Precision Arithmetic

This form of instruction takes operands from the VSRs in double format, performs the operation as if it produced an intermediate result having infinite precision and unbounded exponent range, and then coerces this intermediate result to fit in single-precision format. Status bits, in the FPSCR and optionally in the Condition Register, are set to reflect the single-precision result. The result is then placed into the target VSR in double-precision format. The result lies in the range supported by the single format.

If any input value is not representable in single-precision format and either \(O E=1\) or \(U E=1\), the result placed into the target VSR and the setting of status bits in the FPSCR are undefined.

For xsresp or xsrsqrtesp, if the input value is finite and has an unbiased exponent greater than +127 , the input value is interpreted as an Infinity.
6. Store VSX Scalar Single-Precision
stxsspx converts a single-precision value that is in double-precision format to single-precision format and stores that operand into storage. No floating-point exceptions are caused by stxsspx. (The value being stored is effectively assumed to be the result of an instruction of one of the preceding five types.)

When the result of a Load VSX Scalar Single-Precision (Ixsspx), a VSX Scalar Round to Single-Precision (xsrsp), or a VSX Scalar Single-Precision Arithmetic \({ }^{[1]}\) instruction is stored in a VSR, the low-order 29 bits of FRACTI ON are zero.

\section*{Programming Note}

VSX Scalar Round to Single-Precision (xsrsp) is provided to allow value conversion from double-precision to single-precision with appropriate exception checking and rounding. xsrsp should be used to convert double-precision floating-point values to single-precision values prior to storing them into single format storage elements or using them as operands for single-precision arithmetic instructions. Values produced by single-precision load and arithmetic instructions are already single-precision values and can be stored directly into single format storage elements, or used directly as operands for single-precision arithmetic instructions, without preceding the store, or the arithmetic instruction, by an xsrsp.

\section*{Programming Note}

A single-precision value can be used in double-precision scalar arithmetic operations.

Except for xsresp or xsrsqrtesp, any double-precision value can be used in single-precision scalar arithmetic operations when \(O E=0\) and \(U E=0\). When \(O E=1\) or \(U E=1\), or if the instruction is xsresp or xsrsqrtesp, source operands must be respresentable in single-precision format.

Some implementations may execute single-precision arithmetic instructions faster than double-precision arithmetic instructions. Therefore, if double-precision accuracy is not required, single-precision data and instructions should be used.

\section*{Programming Note}

Both single-precision and double-precision forms are provided for most scalar floating-point instructions. Some scalar floating-point instructions are only provided in double-precision form since their operation is identical to the equivalent scalar single-precision operation.

Of the operations for which only a double-precision form of the instruction is provided,
- instructions that return the absolute value, the negative absolute value, or the negated value (xsnabsdp, xsabsdp, xsnegdp) can be used to perform these operations on scalar single-precision operands,
- instructions that perform a comparison (xscmpodp, xscmpudp) can be used to perform these operations on scalar single-precision operands,
- instructions that determine the maximum (xsmaxdp) or minimum (xsmindp) can be used to perform these operations on scalar single-precision operands, and
- instructions that perform an extraction or insertion of the exponent or significand (xscmpexpdp, xsiexpdp, xststdcdp, xststdcsp, xsxexpdp, xsxsigdp) can be used to perform these operations on scalar single-precision operands.

\footnotetext{
1. VSX Scalar Single-Precision Arithmetic instructions:
xsaddsp, xsdivsp, xsmulsp, xsresp, xssubsp, xsmaddasp, xsmaddmsp, xsmsubasp, xsmsubmsp, xsnmaddasp, xsnmaddmsp, xsnmsubasp, xsnmsubmsp
}

\section*{Integer-Valued Operands}

Instructions are provided to round floating-point operands to integer values in floating-point format. To facilitate exchange of data between the floating-point and integer processing, instructions are provided to convert between floating-point double and single-precision format and integer word and doubleword format in a VSR. Computation on integer-valued operands can be performed using arithmetic instructions of the required precision. (The results might not be integer values.) The three groups of instructions provided specifically to support integer-valued operands are described below.
1. Rounding to a floating-point integer

VSX Scalar Round to Double-Precision Integer \({ }^{[1]}\) instructions round a double-precision operand to an integer value in double-precision format. These instructions can also be used for single-precision operands represented in double-precision format.

VSX Vector Round to Double-Precision Integer \({ }^{[2]}\) instructions round each double-precision vector operand element to an integer value in double-precision format.

VSX Vector Round to Single-Precision Integer \({ }^{[3]}\) instructions round each single-precision vector operand element to an integer value in single-precision format.

Except for xsrdpic, xurdpic, and xvrspic, rounding is performed using the rounding mode specified by the opcode. For xsrdpic, xvrdpic, and xvrspic, rounding is performed using the rounding mode specified by RN.

VSX Round to Floating-Point Integer \({ }^{[4]}\) instructions can cause Invalid Operation (VXSNAN) exceptions.
xsrdpic, xvrdpic, and xvrspic can also cause Inexact exception.
1. VSX Scalar Round to Double-Precision Integer instructions: xsrdpi, xsrdpip, xsrdpim, xsrdpiz, xsrdpic
2. VSX Vector Round to Double-Precision Integer instructions: xvrdpi, xvrdpip, xvrdpim, xvrdpiz, xvrdpic
3. VSX Vector Round to Single-Precision Integer instructions: xvrspi, xvrspip, xvrspim, xvrspiz, xvrspic
4. VSX Round to Floating-Point Integer instructions:
xsrdpi, xsrdpip, xsrdpim, xsrdpiz, xsrdpic, xvrdpi, xvrdpip, xvrdpim, xvrdpiz, xvrdpic, xvrspi, xvrspip, xvrspim, xvrspiz, and xvrspic
5. VSX Scalar Double-Precision to Integer Format Conversion instructions: xscvdpsxds, xscvdpsxws, xscvdpuxds, xscvdpuxws
6. VSX Vector Double-Precision to Integer Format Conversion instructions: xvcvdpsxds, xvcvdpsxws, xvcvdpuxds, xvcvdpuxws
7. VSX Vector Single-Precision to Integer Doubleword Format Conversion instructions: xvcvspsxds, xvcvspuxds
8. VSX Vector Single-Precision to Integer Word Format Conversion instructions: xvcvspsxws, xvcvspuxws
9. VSX Scalar Integer Doubleword to Double-Precision Format Conversion instructions: xscvsxddp, xscvuxddp

\section*{382}
instructions converts a 64-bit signed or unsigned integer to a single-precision floating-point value and returns the result in double-precision format.

> VSX Vector Integer Doubleword to Double-Precision Format Conversion \({ }^{[1]}\) instructions converts the 64-bit signed or unsigned integer in each doubleword element in the source vector operand to double-precision floating-point format.

VSX Vector Integer Word to Double-Precision Format Conversion \({ }^{[2]}\) instructions converts the 32-bit signed or unsigned integer in each odd-numbered word element in the source vector operand to double-precision floating-point format.

\section*{VSX Vector Integer Doubleword to} Single-Precision Format Conversion \({ }^{[3]}\) instructions convert the 64-bit signed or unsigned integer in each doubleword element in the source vector operand to single-precision floating-point format.

VSX Vector Integer Word to Single-Precision Format Conversion \({ }^{[4]}\) instructions convert the 32-bit signed or unsigned integer in each word element in the source vector operand to single-precision floating-point format.

Rounding is performed using the rounding mode specificed in RN. Because of the limitations of the source format, only an Inexact exception can be generated.

\subsection*{7.3.2.6 Rounding}

The material in this section applies to operations that have numeric operands (that is, operands that are not infinities or NaNs). Rounding the intermediate result of such an operation can cause an Overflow exception, an Underflow exception, or an Inexact exception. The remainder of this section assumes that the operation causes no exceptions and that the result is numeric. See Section 7.3.2.2, "Value Representation" and Section 7.4, "VSX Floating-Point Exceptions" for the cases not covered here.

The floating-point arithmetic, and rounding and conversion instructions round their intermediate results. With the exception of the estimate instructions, these instructions produce an intermediate result that
can be regarded as having unbounded precision and exponent range. All but two groups of these instructions normalize or denormalize the intermediate result prior to rounding and then place the final result into the target element of the target VSR in either double-precision, single-precision, or quad-precision format.

The scalar round to double-precision integer, vector round to double-precision integer, and convert double-precision to integer instructions with biased exponents ranging from 1022 through 1074 are prepared for rounding by repetitively shifting the significand right one position and incrementing the biased exponent until it reaches a value of 1075. (Intermediate results with biased exponents 1075 or larger are already integers, and with biased exponents 1021 or less round to zero.) After rounding, the final result for round to double-precision integer instructions is normalized and put in double-precision format, and, for the convert double-precision to integer instructions, is converted to a signed or unsigned integer.

The vector round to single-precision integer and vector convert single-precision to integer instructions with biased exponents ranging from 126 through 178 are prepared for rounding by repetitively shifting the significand right one position and incrementing the biased exponent until it reaches a value of 179. (Intermediate results with biased exponents 179 or larger are already integers, and with biased exponents 125 or less round to zero.) After rounding, the final result for vector round to single-precision integer is normalized and put in double-precision format, and for vector convert single-precision to integer is converted to a signed or unsigned integer.

FR and FI generally indicate the results of rounding. Each of the scalar instructions which rounds its intermediate result sets these bits. There are no vector instructions that modify FR and FI. If the fraction is incremented during rounding, FR is set to 1 , otherwise FR is set to 0 . If the result is inexact, Fl is set to 1 , otherwise FI is set to zero. The scalar round to double-precision integer instructions are exceptions to this rule, setting FR and Fl to 0 . The scalar double-precision estimate instructions set FR and FI to undefined values. The remaining scalar floating-point instructions do not alter FR and FI.

\footnotetext{
10. VSX Scalar Integer Doubleword to Single-Precision Format Conversion instructions: xscvsxdsp, xscvuxdsp
1. VSX Vector Integer Doubleword to Double-Precision Format Conversion instructions: xscvsxddp, xscvuxddp
2. VSX Vector Integer Word to Double-Precision Format Conversion instructions: xscvsxwdp, xscvuxwdp
3. VSX Vector Integer Doubleword to Single-Precision Format Conversion instructions: xscvsxdsp, xscvuxdsp
4. VSX Vector Integer Word to Single-Precision Format Conversion instructions: xscvsxwsp, xscvuxwsp
}

Four user-selectable rounding modes are provided through the Floating-Point Rounding Control field in the FPSCR. See Section 7.2.2, "Floating-Point Status and Control Register" on page 369. These are encoded as follows.

\section*{RN Rounding Mode \\ 00 Round to Nearest Even \\ 01 Round towards Zero \\ 10 Round towards + Infinity \\ 11 Round towards -Infinity}

A fifth rounding mode is provided in the round to floating-point integer instructions (Section 7.6.1.7.2 on page 432), Round to Nearest Away.

A sixth rounding mode is provided in the quad-precision floating-point instructions, Round to Odd.

\section*{Programming Note}

Round to Odd rounding mode is useful when the results of a Quad-Precision Arithmetic instruction are required to be rounded to a shorter precision while avoiding a double rounding error. In this case, the rounding mode of the Quad-Precision Arithmetic instruction is overridden as Round To Odd by setting the RO bit in the instruction encoding to 1, then the result of that Quad-Precision Arithmetic instruction can be rounded to the desired shorter precision using the rounding mode specified in RN by following with a VSX Scalar Round Quad-Precision to Double-Extended-Precision for 15-bit exponent range and 64-bit significand precision, VSX Scalar Round Quad-Precision to Double-Precision for 11-bit exponent range and 53-bit significand precision, or VSX Scalar Round Quad-Precision to Single-Precision for 8-bit exponent range and 24-bit significand precision. For example,
\begin{tabular}{lll} 
xsaddqpo & \(T X, A, B\) & ; use Round to Odd override ( \(R O=1\) ) \\
xsrqpxp & \(T d x p, T x\) & ; final QP result rounded to \(\operatorname{DXP}\)
\end{tabular}

To return a quad-precision result rounded to double-precision requires a 3-instruction sequence,
\begin{tabular}{lll} 
xsaddqpo & Tx, A, B & ; use Round to Odd override (RO=1) \\
xscvapdp & Temp, Tx & ; QP result rounded \& converted to DP \\
xscudpap & Tdp, Temp & ; final QP result rounded to DP
\end{tabular}

To return a quad-precision result rounded to single-precision requires a 4-instruction sequence,
```

xsaddqpo Tx, A, B ; use Round to Odd override (RO=1)
xscvapdpo Temp,Tx ; QP result rounded to DP using Round to Odd \& converted to DP format
xsrsp Temp,Temp ; DP result is rounded to SP
xscudpap Tsp,Temp ; final QP result rounded to SP

```

Let \(Z\) be the intermediate arithmetic result or the operand of a convert operation. If \(Z\) can be represented exactly in the target format, the result in all rounding modes is \(Z\) as represented in the target format. If \(Z\) cannot be represented exactly in the target format, let \(Z 1\) and \(Z 2\) bound \(Z\) as the next larger and next smaller numbers representable in the target format. Then Z1 or Z2 can be used to approximate the result in the target format.

Figure 117 shows the relation of \(Z, Z 1\), and \(Z 2\) in this case. The following rules specify the rounding in the four modes.

See Section 7.3.3.1, "VSX Execution Model for IEEE Operations" on page 386 for a detailed explanation of rounding.

Figure 117 also summarizes the rounding actions for floating-point intermediate result for all supported rounding modes.


Round to Nearest Away
Choose \(Z\) if \(Z\) is representable in the target precision.
Otherwise, choose the value that is closer to \(Z\) ( \(Z 1\) or \(Z 2\) ). In case of a tie, choose the one that is furthest away from 0 .

\section*{Round to Nearest Even}

Choose \(Z\) if \(Z\) is representable in the target precision.
Otherwise, choose the value that is closer to \(Z\) ( \(Z 1\) or \(Z 2\) ). In case of a tie, choose the one that is even (least significant bit is 0 ).

\section*{Round to Odd}

Choose \(Z\) if \(Z\) is representable in the target precision.
Otherwise, choose the value (Z1 or \(Z 2\) ) that is odd (least significant bit is 1 ).

\section*{Round toward Zero}

Choose \(Z\) if \(Z\) is representable in the target precision.
Otherwise, choose the smaller in magnitude (Z1 or Z2).

\section*{Round toward +Infinity}

\section*{Round toward - Infinity}

Choose \(Z\) if \(Z\) is representable in the target precision.
Otherwise, choose \(Z 2\).
Figure 117.Selection of Z1 and Z2

\subsection*{7.3.3 VSX Floating-Point Execution Models}

All implementations of this architecture must provide the equivalent of the following execution models to ensure that identical results are obtained.

Special rules are provided in the definition of the computational instructions for the infinities, denormalized numbers and NaNs. The material in the remainder of this section applies to instructions that have numeric operands and a numeric result (that is, operands and result that are not infinities or NaNs ), and that cause no exceptions. See Section 7.3.2.2 and Section 7.3.3 for the cases not covered here.

Although the double-precision format specifies an 11-bit exponent, exponent arithmetic makes use of two additional bits to avoid potential transient overflow and underflow conditions. One extra bit is required when denormalized double-precision numbers are prenormalized. The second bit is required to permit the computation of the adjusted exponent value in the following cases when the corresponding exception enable bit is 1 :
- Underflow during multiplication using a denormalized operand.
- Overflow during division using a denormalized divisor.
- Undeflow during division using denormalized dividend and a large divisor.

The IEEE standard includes 32 -bit and 64 -bit arithmetic. The standard requires that single-precision arithmetic be provided for single-precision operands.

VSX defines both scalar and vector double-precision floating-point operations to operate only on double-precision operands. VSX also defines vector single-precision floating-point operations to operate only on single-precision operands.

\subsection*{7.3.3.1 VSX Execution Model for IEEE Operations}

IEEE-conforming significand arithmetic is considered to be performed with a floating-point accumulator having the following format, where bits \(0: p \cdot 1\) comprise the significand of the intermediate result (where \(p\) is the length of the significand).


Figure 118.IEEE quad-precision (binary128) floating-point execution model \((p=113)\)


Figure 119.IEEE double-extended-precision
floating-point execution model ( \(p=64\) )


Figure 120.IEEE double-precision (binary64) floating-point execution model ( \(p=53\) )


Figure 121.IEEE single-precision (binary32) floating-point execution model ( \(p=24\) )

The \(S\) bit is the sign bit.
The C bit is the carry bit, which captures the carry out of the significand.

The \(L\) bit is the leading unit bit of the significand, which receives the implicit bit from the operand.

For the quad-precision execution model, FRACTION is a 112-bit field that accepts the fraction of the operand.

For the double-extended-precision execution model, FRACTION is a 63-bit field that accepts the fraction of the operand. This model is used only by the VSX Scalar Round to Double-Extended-Precision instruction.

For the double-precision execution model, FRACTION is a 52-bit field that accepts the fraction of the operand.

For the single-precision execution model, FRACTI ON is a 23-bit field that accepts the fraction of the operand.

The Guard ( \(G\) ), Round (R), and Sticky ( \(X\) ) bits are extensions to the low-order bits of the accumulator to provide the effect of an unbounded significand. The G and \(R\) bits are required for postnormalization of the result. The \(G, R\), and \(X\) bits are required during rounding to determine if the intermediate result is equally near the two nearest representable values. The \(X\) bit serves as an extension to the \(G\) and \(R\) bits by representing the logical OR of all bits that appear to the low-order side of the R bit, resulting from either shifting the accumulator right or to other generation of low-order result bits. The \(G\) and \(R\) bits participate in the left shifts with zeros being shifted into the \(R\) bit. Table 4 shows the significance of the \(G, R\), and \(X\) bits with respect to the intermediate result (IR), the representable number next lower in magnitude (NL),
and the representable number next higher in magnitude (NH).
\begin{tabular}{c|c|c|l}
\hline \(\mathbf{G}\) & \(\mathbf{R}\) & \(\mathbf{X}\) & \multicolumn{1}{|c}{ Interpretation } \\
\hline 0 & 0 & 0 & IR is exact \\
\hline 0 & 0 & 1 & IR closer to NL \\
\hline 0 & 1 & 0 & \\
\hline 0 & 1 & 1 & \\
\hline 1 & 0 & 0 & IR midway between NL and NH \\
\hline 1 & 0 & 1 & IR closer to NH \\
\hline 1 & 1 & 0 & \\
\hline 1 & 1 & 1 & \\
\hline
\end{tabular}

Table 4. Interpretation of G, R, and \(X\) bits
Table 5 shows the positions of the Guard, Round, and Sticky bits for double-precision and single-precision floating-point numbers relative to the accumulator illustrated in Figures 112, 113, 114, and 115.
\begin{tabular}{c|c|c|c} 
Format & Guard & Round & Sticky \\
\hline Double & G bit & R bit & X bit \\
\hline Single & 24 & 25 & OR of bits \(26: 52, G, R, X\) \\
\hline
\end{tabular}

Table 5. Location of the Guard, Round, and Sticky bits in the IEEE execution model

The significand of the intermediate result is prepared for rounding by shifting its contents right, if required, until the least significant bit to be retained is in the low-order bit position of the fraction.

I Six rounding modes are provided as described in Section 7.3.2.6, "Rounding" on page 383. The rules for rounding in each mode are as follows.

\section*{- Round to Nearest Even}

If \(\mid R\) is exact, choose I \(R\). Otherwise, if | \(R\) is closer to \(N L\), choose \(N L\). Otherwise, if I \(R\) is closer to \(N H\), choose \(N H\). Otherwise, if \(\mid R\) is midway between \(N L\) and \(N H\), choose whichever of NL and NH is even.

\section*{- Round towards Zero}

If \(\mid R\) is exact, choose \(\mid R\).
Otherwise, choose NL.

\section*{- Round towards +Infinity}

If \(\mid R\) is exact, choose \(\mid R\).
Otherwise, if positive, choose NH.
Otherwise, if negative, choose NL .

\section*{- Round towards -Infinity}

If \(\mid R\) is exact, choose \(\mid R\).
Otherwise, if positive, choose NL.
Otherwise, if negative, choose NH.
- Round to Nearest Away

If \(\mid R\) is exact, choose \(\mid R\).
Otherwise, if \(\mathrm{G}=0\), choose NL.
Otherwise, if \(\mathrm{G}=1\), choose NH .
- Round to Odd

If \(\mid R\) is exact, choose \(\mid R\).
Otherwise, choose \(N L\), and if \(G=1, R=1\), or \(X=1\), the least-significant bit of the result is set to 1 .

Four of the rounding modes are user-selectable through RN.
\begin{tabular}{ll} 
RN & Rounding Mode \\
Ob00 & Round to Nearest Even \\
Ob01 & Round toward Zero \\
Ob10 & Round toward + Infinity \\
Ob11 & Round toward - Infinity
\end{tabular}

Round to Nearest Away is provided in the VSX Round to Floating-Point Integer instructions (Section 7.6.1.7.2 on page 432).

Round to Odd is provided in the VSX Quad-Precision Floating-Point Arithmetic instructions as an override to the rounding mode selected by RN with the rules for rounding as follows.

If \(G=1, R=1\), or \(X=1\), the result is inexact.
If rounding results in a carry into C , the significand is shifted right one position and the exponent is incremented by one. This yields an inexact result, and possibly also exponent overflow. Fraction bits are stored to the target VSR.

\subsection*{7.3.3.2 VSX Execution Model for Multiply-Add Type Instructions}

This architecture provides a special form of instruction that performs up to three operations in one instruction (a multiplication, an addition, and a negation). With this added capability comes the special ability to produce a more exact intermediate result as input to the rounder. 32-bit arithmetic is similar, except that the FRACTION field is smaller.

Multiply-add significand arithmetic is considered to be performed with a floating-point accumulator having the

\section*{Version 3.0}
following format, where bits 0:106 comprise the significand of the intermediate result.
\begin{tabular}{c|c|c|c|c|}
\hline\(S\) & C & L & FRACTION & X \\
\hline 0 & 1 & 2 & 3 & \\
\hline
\end{tabular}

Figure 122.Multiply-add 64-bit execution model
The first part of the operation is a multiplication. The multiplication has two 53-bit significands as inputs, which are assumed to be prenormalized, and produces a result conforming to the above model. If there is a carry out of the significand (into the \(C\) bit), the significand is shifted right one position, shifting the \(L\) bit (leading unit bit) into the most significant bit of the FRACTI ON and shifting the C bit (carry out) into the \(L\) bit. All 106 bits ( L bit, the FRACTI ON) of the product take part in the add operation. If the exponents of the two inputs to the adder are not equal, the significand of the operand with the smaller exponent is aligned (shifted) to the right by an amount that is added to that exponent to make it equal to the other input's exponent. Zeros are shifted into the left of the significand as it is aligned and bits shifted out of bit 105 of the significand are ORed into the \(X^{\prime}\) bit. The add operation also produces a result conforming to the above model with the \(X^{\prime}\) bit taking part in the add operation.

The result of the addition is then normalized, with all bits of the addition result, except the \(X^{\prime}\) bit, participating in the shift. The normalized result serves as the intermediate result that is input to the rounder.

For rounding, the conceptual Guard, Round, and Sticky bits are defined in terms of accumulator bits. Figure 6 shows the positions of the Guard, Round, and Sticky bits for double-precision and single-precision floating-point numbers in the multiply-add execution model.
\begin{tabular}{c|c|c|c}
\hline Format & Guard & Round & Sticky \\
\hline Double & 53 & 54 & OR of 55:105, X' \\
\hline Single & 24 & 25 & OR of \(26: 105, X^{\prime}\) \\
\hline
\end{tabular}

Table 6. Location of the Guard, Round, and Sticky bits in the multiply-add execution model

The rules for rounding the intermediate result are the same as those given in Section 7.3.3.1.

If the instruction is a negative multiply-add or negative multiply-subtract type instruction, the final result is negated.

\subsection*{7.4 VSX Floating-Point Exceptions}

This architecture defines the following floating-point exceptions under the IEEE-754 exception model:
- Invalid Operation exception

\author{
SNaN \\ Infinity-Infinity \\ Infinity-Infinity \\ Zero:Zero \\ Infinity×Zero \\ Invalid Compare \\ Software-Defined Condition \\ Invalid Square Root \\ Invalid Integer Convert
}
- Zero Divide exception
- Overflow exception
- Underflow exception
- Inexact exception

These exceptions, other than Invalid Operation exception resulting from a Software-Defined Condition, can occur during execution of computational instructions. An Invalid Operation exception resulting from a Software-Defined Condition occurs when a Move To FPSCR instruction sets VXSOFT to 1.

Each floating-point exception, and each category of Invalid Operation exception, has an exception bit in the FPSCR. In addition, each floating-point exception has a corresponding enable bit in the FPSCR. The exception bit indicates the occurrence of the corresponding exception. If an exception occurs, the corresponding enable bit governs the result produced by the instruction and, in conjunction with the FEO and FEl bits (see page 390), whether and how the system floating-point enabled exception error handler is invoked. In general, the enabling specified by the enable bit is of invoking the system error handler, not of permitting the exception to occur. The occurrence of an exception depends only on the instruction and its inputs, not on the setting of any control bits. The only deviation from this general rule is that the occurrence of an Underflow exception depends on the setting of the enable bit.

A single instruction, other than mtfsfi or mtfsf, can set more than one exception bit only in the following cases:
- An Inexact exception can be set with an Overflow exception.
- An Inexact exception can be set with an Underflow exception.
- An Invalid Operation exception ( SNaN ) is set with an Invalid Operation exception (Infinity \(\times 0\) ) for multiply-add class instructions for which the values being multiplied are infinity and zero and the value being added is an SNaN .
- An Invalid Operation exception (SNaN) can be set with an Invalid Operation exception (Invalid Compare) for ordered comparison instructions.
- An Invalid Operation exception (SNaN) can be set with an Invalid Operation exception (Invalid Integer Convert) for convert to integer instructions.

When an exception occurs, the writing of a result to the target register can be suppressed, or a result can be delivered, depending on the exception.

The writing of a result to the target register is suppressed for the certain kinds of exceptions, based on whether the instruction is a vector or a scalar instruction, so that there is no possibility that one of the operands is lost. For other kinds of exceptions and also depending on whether the instruction is a vector or a scalar instruction, a result is generated and written to the destination specified by the instruction causing the exception. The result can be a different value for the enabled and disabled conditions for some of these exceptions. Table 7 lists the types of exceptions and indicates whether a result is written to the target VSR or suppressed.
\begin{tabular}{l|c|c}
\hline \multicolumn{1}{c|}{ On exception type... } & \begin{tabular}{c} 
Scalar \\
Instruction \\
Results
\end{tabular} & \begin{tabular}{c} 
Vector \\
Instruction \\
Results
\end{tabular} \\
\hline Enabled Invalid Operation & suppressed & suppressed \\
\hline Enabled Zero Divide & suppressed & suppressed \\
\hline Enabled Overflow & written & suppressed \\
\hline Enabled Underflow & written & suppressed \\
\hline Enabled Inexact & written & suppressed \\
\hline Disabled Invalid Operation & written & written \\
\hline
\end{tabular}

Table 7. Exception Types Result Suppression
\begin{tabular}{l|c|c}
\hline \multicolumn{1}{c|}{ On exception type... } & \begin{tabular}{c} 
Scalar \\
Instruction \\
Results
\end{tabular} & \begin{tabular}{c} 
Vector \\
Instruction \\
Results
\end{tabular} \\
\hline Disabled Zero Divide & written & written \\
\hline Disabled Overflow & written & written \\
\hline Disabled Underflow & written & written \\
\hline Disabled Inexact & written & written \\
\hline
\end{tabular}

Table 7. Exception Types Result Suppression
The subsequent sections define each of the floating-point exceptions and specify the action that is taken when they are detected.

The IEEE standard specifies the handling of exceptional conditions in terms of traps and trap handlers. In this architecture, an FPSCR exception enable bit of 1 causes generation of the result value specified in the IEEE standard for the trap enabled case; the expectation is that the exception is detected by software, which revises the result. An FPSCR exception enable bit of 0 causes generation of the default result value specified for the trap disabled (or no trap occurs or trap is not implemented) case. The expectation is that the exception is not detected by software, which uses the default result. The result to be delivered in each case for each exception is described in the following sections.

The IEEE default behavior when an exception occurs is to generate a default value and not to notify software. In this architecture, if the IEEE default behavior when an exception occurs is required for all exceptions, all FPSCR exception enable bits must be set to 0 , and Ignore Exceptions Mode (see below) should be used. In this case, the system floating-point enabled exception error handler is not invoked, even if floating-point exceptions occur: software can inspect the FPSCR exception bits, if necessary, to determine whether exceptions have occurred.

In this architecture, if software is to be notified that a given kind of exception has occurred, the corresponding FPSCR exception enable bit must be set to 1, and a mode other than Ignore Exceptions Mode must be used. In this case, the system floating-point enabled exception error handler is invoked if an enabled floating-point exception occurs. The system floating-point enabled exception error handler is also invoked if a Move To FPSCR instruction causes an exception bit and the corresponding enable bit both to be 1. The Move To FPSCR instruction is considered to cause the enabled exception.

The FEO and FE1 bits control whether and how the system floating-point enabled exception error handler is invoked if an enabled floating-point exception occurs. The location of these bits and the requirements
for altering them are described in Book III. The system floating-point enabled exception error handler is never invoked because of a disabled floating-point exception. The effects of the four possible settings of these bits are as follows.

\section*{FEO FE1 Description}

00 Ignore Exceptions Mode
Floating-point exceptions do not cause the system floating-point enabled exception error handler to be invoked.

01 Imprecise Nonrecoverable Mode
The system floating-point enabled exception error handler is invoked at some point at or beyond the instruction that caused the enabled exception. It may not be possible to identify the excepting instruction or the data that caused the exception. Results produced by the excepting instruction might have been used by or might have affected subsequent instructions that are executed before the error handler is invoked.

10 Imprecise Recoverable Mode
The system floating-point enabled exception error handler is invoked at some point at or beyond the instruction that caused the enabled exception. Sufficient information is provided to the error handler for it to identify the excepting instruction, the operands, and correct the result. No results produced by the excepting instruction have been used by or affected subsequent instructions that are executed before the error handler is invoked.

11 Precise Mode
The system floating-point enabled exception error handler is invoked precisely at the instruction that caused the enabled exception.

In all cases, the question of whether a floating-point result is stored, and what value is stored, is governed by the FPSCR exception enable bits, as described in subsequent sections, and is not affected by the value of the FEO and FE1 bits.

In all cases in which the system floating-point enabled exception error handler is invoked, all instructions before the instruction at which the system floating-point enabled exception error handler is invoked have been completed, and no instruction after the instruction at which the system floating-point enabled exception error handler is invoked has begun execution. The instruction at which the system floating-point enabled exception error handler is invoked has completed if it is the excepting instruction,
and there is only one such instruction. Otherwise, it has not begun execution, or has been partially executed in some cases, as described in Book III.

\section*{Programming Note}

In any of the three non-Precise modes, a Floating-Point Status and Control Register instruction can be used to force any exceptions, because of instructions initiated before the Floating-Point Status and Control Register instruction, to be recorded in the FPSCR. (This forcing is superfluous for Precise Mode.)

In both Imprecise modes, a Floating-Point Status and Control Register instruction can be used to force any invocations of the system floating-point enabled exception error handler that result from instructions initiated before the Floating-Point Status and Control Register instruction to occur. This forcing has no effect in Ignore Exceptions Mode, and is superfluous for Precise Mode.

The last sentence of the paragraph preceding this Programming Note can apply only in the Imprecise modes, or if the mode has just been changed from Ignore Exceptions Mode to some other mode. It always applies in the latter case.

To obtain the best performance across the widest range of implementations, the programmer should obey the following guidelines.
- If the IEEE default results are acceptable to the application, Ignore Exceptions Mode should be used with all FPSCR exception enable bits set to 0.
- If the IEEE default results are not acceptable to the application, Imprecise Nonrecoverable Mode should be used, or Imprecise Recoverable Mode if recoverability is needed, with FPSCR exception enable bits set to 1 for those exceptions for which the system floating-point enabled exception error handler is to be invoked.
- Ignore Exceptions Mode should not, in general, be used when any FPSCR exception enable bits are set to 1 .
- Precise Mode can degrade performance in some implementations, perhaps substantially, and therefore should be used only for debugging and other specialized applications.

\subsection*{7.4.1 Floating-Point Invalid Operation Exception}

\subsection*{7.4.1.1 Definition}

An Invalid Operation exception occurs when an operand is invalid for the specified operation. The invalid operations are:

\section*{SNaN}

Any floating-point operation on a Signaling NaN.

\section*{Infinity-Infinity}

Magnitude subtraction of infinities.

\section*{Infinity-Infinity}

Floating-point division of infinity by infinity.

\section*{Zero \(\div\) Zero}

Floating-point division of zero by zero.

\section*{Infinity \(\times\) Zero}

Floating-point multiplication of infinity by zero.

\section*{Invalid Compare}

Floating-point ordered comparison involving a NaN .

\section*{Invalid Square Root}

Floating-point square root or reciprocal square root of a nonzero negative number.

\section*{Invalid Integer Convert}

Floating-point-to-integer convert involving a number too large in magnitude to be represented in the target format, or involving an infinity or a NaN .

An Invalid Operation exception also occurs when an \(\boldsymbol{m t f s f i}, \boldsymbol{m t f s f}\), or mtfsb1 instruction is executed that sets VXSOFT to 1 (Software-Defined Condition).

The action to be taken depends on the setting of the Invalid Operation Exception Enable bit of the FPSCR.

\subsection*{7.4.1.2 Action for VE=1}

When Invalid Operation exception is enabled \((V E=1)\) and an Invalid Operation exception occurs, the following actions are taken:

For VSX Scalar Floating-Point Arithmetic, VSX Scalar DP-SP Conversion, VSX Scalar Convert Floating-Point to Integer, and VSX Scalar Round to Floating-Point Integer instructions:
1. One or two of the following Invalid Operation exceptions are set to 1.
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXI SI & (if Infinity-Infinity) \\
VXI DI & (if Infinity \(\div\) Infinity) \\
VXZDZ & (if Zero \(\div\) Zero) \\
VXI MZ & (if Infinity×Zero) \\
VXSQRT & (if Invalid Square Root) \\
VXCVI & (if Invalid Integer Convert)
\end{tabular}
2. Update of VSR[ XT] is suppressed.
3. \(F R\) and \(F l\) are set to zero.
4. \(F P R F\) is unchanged.

For VSX Scalar Floating-Point Compare instructions:
1. One or two of the following Invalid Operation exceptions are set to 1.
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXVC & (if Invalid Compare)
\end{tabular}
2. \(F R, F I\), and \(C\) are unchanged.
3. FPCC is set to reflect unordered.

For any of the following instructions,
```

VSX Scalar Quad-Precision Arithmetic instructions:
xsaddqp[o], xsdivqp[o], xsmulqp[o], xssqrtqp[o], xssubqp[o]
xsmaddqp[o], xsmsubqp[o], xsnmaddqp[o], xsnmsubqp[o]
VSX Scalar Quad-Precision Convert to Integer instructions:
xscvqpsdz, xscvqpswz, xscvqpudz, xscvqpuwz

```
VSX Scalar Round Quad-Precision to Double-Extended-Precision (xsrqpxp)
VSX Scalar Round to Quad-Precision Integer (xsrqpi)
VSX Scalar Convert Quad-Precision to Double-Precision [using round to Odd] (xscvqpdp[o])
do the following.
1. One or two of the following Invalid Operation exceptions are set to 1.
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXI SI & (if Infinity - Infinity) \\
VXI DI & (if Infinity \(\div\) Infinity) \\
VXZDZ & (if Zero \(\div\) Zero) \\
VXI MZ & (if Infinity \(\times\) Zero) \\
VXSQRT & (if Invalid Square Root) \\
VXCVI & (if Invalid Integer Convert)
\end{tabular}
2. VSR[VRT+32] is not modified.
3. \(F R\) and \(F I\) are set to zero. \(F P R F\) is not modified.

For any of the following instructions,
VSX Scalar Compare Ordered Quad-Precision (xscmpoqp)
VSX Scalar Compare Unordered Quad-Precision (xscmpuqp)
do the following.
1. One or two of the following Invalid Operation exceptions are set to 1.
\begin{tabular}{ll} 
VXSNAN & (if SNaN ) \\
VXVC & (if Invalid Compare)
\end{tabular}
2. \(F R, F I\), and \(C\) are not modified. \(F P C C\) is set to reflect unordered.

For any of the following instructions,
VSX Scalar Convert Half-Precision to Double-Precision (xscvdphp)
VSX Scalar Convert Double-Precision to Half-Precision with round (xscvdphp)
do the following.
1. VXSNAN is set to 1 .
2. VSR[ XT] is not modified.
3. \(F R\) and Fl are set to \(0 . \mathrm{FPRF}\) is not modified.

For any of the following instructions,

\section*{Version 3.0}
do the following.
1. VXSNAN is set to 1 .
2. VSR[XT] is not modified.
3. \(F R, F I\), and \(F P R F\) are not modified.

For any of the following instructions,
VSX Vector Floating-Point Arithmetic instructions:
VSX Vector Floating-Point Compare instructions:
VSX Vector DP-SP Conversion instructions:
VSX Vector Convert Floating-Point to Integer instructions:
VSX Vector Round to Floating-Point Integer instructions:
do the following.
1. One or two of the following Invalid Operation exceptions are set to 1 .
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXI SI & (if Infinity - Infinity) \\
VXI DI & (if Infinity \(\div\) Infinity) \\
VXZDZ & (if Zero \(\div\) Zero) \\
VXI MZ & (if Infinity \(\times\) Zero) \\
VXVC & (if Invalid Compare) \\
VXSQRT & (if Invalid Square Root) \\
VXCVI & (if Invalid Integer Convert)
\end{tabular}
2. Update of VSR[XT] is suppressed for all vector elements.
3. \(F R\) and \(F I\) are unchanged.
4. \(F P R F\) is unchanged.

\subsection*{7.4.1.3 Action for VE=0}

When Invalid Operation exception is disabled \((V E=0)\) and an Invalid Operation exception occurs, the following actions are taken:

For the VSX Scalar round and Convert Double-Precision to Single-Precision format (xscvdpsp) instruction:
1. VXSNAN is set to 1 .
2. The single-precision representation of a Quiet NaN is placed into word element 0 of VSR[XT]. The contents of word elements 1-3 of VSR[ XT] are undefined.
3. FR and Fl are set to 0 .
4. \(\quad \mathrm{FPRF}\) is set to indicate the class of the result (Quiet NaN ).

For the VSX Vector Single-Precision Arithmetic instructions, VSX Vector Single-Precision Maximum/Minimum instructions, the VSX Vector round and Convert Double-Precision to Single-Precision format (xvcvdpsp) instruction, and the VSX Vector Round to Single-Precision Integer instructions:
1. One or two of the following Invalid Operation exceptions are set to 1.
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXI SI & (if Infinity - Infinity) \\
VXI DI & (if Infinity \(\div\) Infinity) \\
VXZDZ & (if Zero \(\div\) Zero) \\
VXI MZ & (if Infinity \(\times\) Zero) \\
VXSQRT & (if Invalid Square Root)
\end{tabular}
2. The single-precision representation of a Quiet NaN is placed into its respective word element of VSR[ XT].
3. \(F R, F I\), and \(F P R F\) are not modified.

For the VSX Scalar Double-Precision Arithmetic instructions, VSX Scalar Double-Precision Maximum/Minimum instructions, the VSX Scalar Convert Single-Precision to Double-Precision format (xscvspdp) instruction, and the VSX Scalar Round to Double-Precision Integer instructions:
1. One or two of the following Invalid Operation exceptions are set to 1 .
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXI SI & (if Infinity - Infinity) \\
VXI DI & (if Infinity \(\div\) Infinity) \\
VXZDZ & (if Zero \(\div\) Zero) \\
VXI MZ & (if Infinity \(\times\) Zero) \\
VXSQRT & (if Invalid Square Root)
\end{tabular}
2. The double-precision representation of a Quiet NaN is placed into doubleword element 0 of VSR[ XT] . The contents of doubleword element 1 of VSR[ XT] are undefined.
3. \(F R\) and \(F I\) are set to 0 .
4. \(\quad \mathrm{FPRF}\) is set to indicate the class of the result (Quiet NaN ).

For any of the following instructions,
VSX Scalar Quad-Precision Arithmetic instructions:
xsaddqp[o], xsdivqp[o], xsmulqp[o], xssqrtqp[o], xssubqp[o]
xsmaddqp[o], xsmsubqp[o], xsnmaddqp[o], xsnmsubqp[o]
VSX Scalar Quad-Precision Round to Integer (xsrqpi)
do the following.
1. One or two of the following Invalid Operation exceptions are set to 1.
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXI SI & (if Infinity \(\cdot\) Infinity) \\
VXI DI & (if Infinity \(\div\) Infinity) \\
VXZDZ & (if Zero \(\div\) Zero) \\
VXI MZ & (if Infinity \(\times\) Zero) \\
VXSQRT & (if Invalid Square Root)
\end{tabular}
2. The quad-precision representation of a Quiet NaN is placed into VSR[ VRT + 32].
3. \(F R\) and \(F I\) are set to 0 . \(F P R F\) is set to indicate the class of the result (Quiet NaN ).

For VSX Scalar Round Quad-Precision to Double-Extended-Precision (xsrqpxp), do the following.
1. VXSNAN is set to 1 .
2. The Quiet NaN is placed into VSR[ VRT +32] in quad-precision format.
3. \(F R\) and \(F I\) are set to 0 . FPRF is set to indicate the class of the result (Quiet NaN ).

For any of the following instructions,
VSX Scalar Compare Ordered Quad-Precision (xscmpoqp)
VSX Scalar Compare Unordered Quad-Precision (xscmpoqp)
do the following.
1. One or two of the following Invalid Operation exceptions are set to 1 .
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXVC & (if Invalid Compare)
\end{tabular}
2. \(F R, F I\) and \(C\) are unchanged. \(F P C C\) is set to reflect unordered.

For VSX Scalar Convert Quad-Precision to Double-Precision [using round to Odd] (xscvqpdp[o]), do the following.
1. VXSNAN is set to 1 .
2. The double-precision Quiet NaN result is placed into doubleword element 0 of VSR[VRT+32] in double-precision format.
\(0 \times 0000_{-} 0000^{2} 0000_{-} 0000\) is placed into doubleword element 1 of VSR[VRT+32].
3. \(F R\) and \(F I\) are set to 0 . \(F P R F\) is set to indicate the class of the result (Quiet NaN ).

For VSX Scalar Convert Quad-Precision to Signed Doubleword (xscvqpsdz), do the following.
1. One or two of the following Invalid Operation exceptions are set to 1.

VXSNAN (if SNaN)
VXCVI (if Invalid Integer Convert)
2. OXTFF_FFFF_FFF_FFFF is placed into doubleword element 0 of VSR[VRT+32] if the quad-precision operand in VSR[VRB+32] is a positive number or + Infinity.
\(0 \times 8000 \_0000 \_0000 \_0000\) is placed into doubleword element 0 of VSR[VRT +32 ] if the quad-precision operand in VSR[ VRB +32 ] is a negative number, - Infinity, or NaN.
\(0 \times 0000 \_0000 \_0000 \_0000\) is placed into doubleword element 1 of VSR[VRT +32 ].
3. \(F R\) and \(F I\) are set to \(0 . F P R F\) is undefined.

For VSX Scalar Convert Quad-Precision to Signed Word (xscvqpswz), do the following.
1. One or two of the following Invalid Operation exceptions are set to 1.

VXSNAN (if SNaN)
VXCVI (if Invalid Integer Convert)
2. \(\quad 0 \times 7 F F F\) FFFF is placed into word element 1 of VSR[VRT+32] if the quad-precision operand in VSR[VRB +32 ] is a positive number or +Infinity.
\(0 \times 8000 \_0000\) is placed into word element 1 of VSR[VRT+32] if the quad-precision operand in VSR[VRB+32] is a negative number, - Infinity, or NaN .
\(0 \times 0000 \_0000\) is placed into word elements 0,2 , and 3 of VSR[VRT+32].
3. \(F R\) and \(F I\) are set to \(0 . F P R F\) is undefined.

For VSX Scalar Convert Quad-Precision to Unsigned Doubleword (xscvqpudz), do the following.
1. One or two of the following Invalid Operation exceptions are set to 1.

VXSNAN (if SNaN )
VXCVI (if Invalid Integer Convert)
2. OXFFF_FFFFFFFF_FFFF is placed into doubleword element 0 of VSR[VRT+32] if the quad-precision operand in V'SR[ VRB +32\(]\) is a positive number or + Infinity.
\(0 \times 00000_{-} 0000 \_0000 \_0000\) is placed into doubleword element 0 of VSR[VRT+32] if the quad-precision operand in VSR[VRB+32] is a negative number, - Infinity, or NaN .
\(0 \times 0000_{-} 00000^{0} 0000_{-} 0000\) is placed into doubleword element 1 of VSR[VRT+32].
3. \(F R\) and \(F I\) are set to \(0 . F P R F\) is undefined.

For VSX Scalar Convert Quad-Precision to Unsigned Word (xscvqpuwz), do the following.
1. One or two of the following Invalid Operation exceptions are set to 1.

VXSNAN (if SNaN)
VXCVI (if Invalid Integer Convert)
2. OXFFFF_FFFF is placed into word element 1 of VSR[VRT+32] if the quad-precision operand in VSR[VRB+32] is a positive number or +Infinity.
\(0 \times 0000 \_0000\) is placed into word element 1 of VSR[VRT+32] if the quad-precision operand in VSR[VRB +32 ] is a negative number, - Infinity, or NaN .
\(0 \times 0000 \_0000\) is placed into word elements 0,2 , and 3 of VSR[VRT +32\(]\).
3. \(F R\) and \(F I\) are set to \(0 . F P R F\) is undefined.

For VSX Scalar Convert Double-Precision to Half-Precision with round (xscvdphp), do the following.
1. VXSNAN is set to 1 .
2. The half-precision representation of a Quiet NaN is placed into the rightmost halfword of doubleword element 0 of VSR[ XT]. The contents of the leftmost 3 halfwords of doubleword element 0 of VSR[XT] are set to 0 . The contents of doubleword element 1 of VSR[ XT] are undefined.
3. FR and FI are set to \(0 . \mathrm{FPRF}\) is set to indicate the class of the result (Quiet NaN ).

For VSX Scalar Convert Half-Precision to Double-Precision (xscvhpdp), do the following.
1. VXSNAN is set to 1 .
2. The double-precision representation of a Quiet NaN is placed into doubleword element 0 of VSR[ XT ]. The contents of doubleword element 1 of VSR[ XT] are undefined.
3. \(F R\) and Fl are set to 0 . \(F P R F\) is set to indicate the class of the result (Quiet NaN ).

For the VSX Vector Double-Precision Arithmetic instructions, VSX Vector Double-Precision Maximum/Minimum instructions, the VSX Vector Convert Single-Precision to Double-Precision format (xvcvspdp) instruction, and the VSX Vector Round to Double-Precision Integer instructions:
1. One or two of the following Invalid Operation exceptions are set to 1 .

VXSNAN (if SNaN)
VXI SI (if Infinity - Infinity)
VXI DI (if Infinity \(\div\) Infinity)
VXZDZ (if Zero \(\div\) Zero)
VXI MZ (if Infinity \(\times\) Zero)
VXSQRT (if Invalid Square Root)
2. The double-precision representation of a Quiet NaN is placed into its respective doubleword element of VSR[ XT] .
3. \(F R, F I\), and \(F P R F\) are not modified.

For the VSX Scalar Convert Double-Precision to Signed Integer Doubleword (xscvdpsxd) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1.
\(\begin{array}{ll}\text { VXSNAN } & \text { (if SNaN) } \\ \text { VXCVI } & \text { (if Invalid Integer Convert) }\end{array}\)
2. OXTFFF_FFFF_FFF_FFFF is placed into doubleword element 0 of VSR[XT] if the double-precision operand in doubleword element 0 of VSR[ XB] is a positive number or + Infinity.
\(0 \times 80000^{2} 0000_{-} 0000_{-} 0000\) is placed into doubleword element 0 of VSR[XT] if the double-precision operand in doubleword element 0 of VSR[ XB] is a negative number, - Infinity, or NaN.

The contents of doubleword element 1 of VSR[ XT] are undefined.
3. \(F R\) and \(F I\) are set to 0 .
4. \(F P R F\) is undefined.

For the VSX Scalar Convert Double-Precision to Unsigned Integer Doubleword (xscvdpuxd) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1 .
\(\begin{array}{ll}\text { VXSNAN } & \text { (if } \mathrm{SNaN} \text { ) } \\ \text { VXCVI } & \text { (if Invalid Integer Convert) }\end{array}\)
 in doubleword element 0 of VSR[ XB] is a positive number or + Infinity.
\(0 \times 0000 \_0000 \_0000 \_0000\) is placed into doubleword element 0 of VSR[XT] if the double-precision operand in doubleword element 0 of VSR[ XB] is a negative number, - Infinity, or NaN .

The contents of doubleword element 1 of VSR[ XT] are undefined.
3. FR and FI are set to 0 .
4. FPRF is undefined.

For the VSX Scalar Convert Double-Precision to Signed Integer Word (xscvdpsxw) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1.

VXSNAN (if SNaN)

VXCVI (if Invalid Integer Convert)
2. \(0 \times 7 F F F \_F F F F\) is placed into word element 1 of VSR[XT] if the double-precision operand in doubleword element 0 of VSR[ XB] is a positive number or + Infinity.

0x8000_0000 is placed into word element 1 of VSR[ XT] if the double-precision operand in doubleword element 0 of VSR[ \(X B]\) is a negative number, - Infinity, or NaN .

The contents of word elements 0,2 , and 3 of VSR[ XT] are undefined.
3. \(F R\) and \(F I\) are set to 0 .
4. \(F P R F\) is undefined.

For the VSX Scalar Convert Double-Precision to Unsigned Integer Word (xscvdpuxw) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1.

VXSNAN (if SNaN)
VXCVI (if Invalid Integer Convert)
2. \(0 X F F F F\) _FFFF is placed into word element 1 of VSR[XT] if the double-precision operand in doubleword element 0 of VSR[ \(X B]\) is a positive number or +Infinity.
\(0 \times 0000 \_0000\) is placed into word element 1 of VSR[ XT] if the double-precision operand in doubleword element 0 of VSR[XB] is a negative number, - Infinity, or NaN .

The contents of word elements 0,2 , and 3 of VSR[ XT] are undefined.
3. \(F R\) and \(F I\) are set to 0 .
4. \(F P R F\) is undefined.

For the VSX Vector Convert Double-Precision to Signed Integer Doubleword (xvcvdpsxd) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1.

VXSNAN (if SNaN)
VXCVI (if Invalid Integer Convert)
2. \(0 \times 7 F F F \_\)FFFF_FFFF_FFFF is placed into doubleword element \(i\) of VSR[XT] if the double-precision operand in the corresponding doubleword element of \(\operatorname{VSR}[X B]\) is a positive number or +Infinity.

0x8000_0000_0000_0000 is placed into its respective doubleword element i of VSR[XT] if the double-precision operand in the corresponding doubleword element of \(\operatorname{VSR}[X B]\) is a negative number, - Infinity, or NaN.
3. \(\mathrm{FR}, \mathrm{FI}\), and FPRF are not modified.

For the VSX Vector Convert Double-Precision to Unsigned Integer Doubleword (xvcvdpuxd) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1.

VXSNAN (if SNaN)
VXCVI (if Invalid Integer Convert)
 in doubleword element \(i\) of \(\operatorname{VSR}[X B]\) is a positive number or +Infinity.
\(0 \times 0000 \_0000 \_0000 \_0000\) is placed into doubleword element \(i\) of VSR[ XT] if the double-precision operand in doubleword element i of VSR[ XB] is a negative number, - Infinity, or NaN.
3. \(F R, F I\), and \(F P R F\) are not modified.

For the VSX Vector Convert Double-Precision to Signed Integer Word (xvcvdpsxw) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1.
```

VXSNAN (if SNaN)
VXCVI (if Invalid Integer Convert)

```
2. \(0 \times 7 F F F\) FFFF is placed intoword element \(i \times 2\) of \(\operatorname{VSR}[X T]\) if the double-precision operand in doubleword element \(i\) of VSR[ \(X B]\) is a positive number or + Infinity.
\(0 \times 8000 \_0000\) is placed into word element \(i \times 2\) of VSR[ XT] if the double-precision operand in doubleword element i of \(\mathrm{VSR}[\mathrm{XB}]\) is a negative number, - Infinity, or NaN .

The contents of word element x 2 t 1 of VSR[ XT] are undefined.
3. \(F R, F I\), and \(F P R F\) are not modified.

For the VSX Vector Convert Double-Precision to Unsigned Integer Word (xvcvdpuxw) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1.
\begin{tabular}{ll} 
VXSNAN & (if SNaN ) \\
VXCVI & (if Invalid Integer Convert)
\end{tabular}
2. OXFFFF FFFF is placed into word element \(\mathrm{i} \times 2\) of VSR[XT] if the double-precision operand in doubleword element \(i\) of VSR[ \(X B]\) is a positive number or +Infinity.
\(0 \times 0000 \_0000\) is placed into word element \(\mathrm{i} \times 2\) of VSR[ XT] if the double-precision operand in doubleword element i of VSR[XB] is a negative number, - Infinity, or NaN .

The contents of word element \(\mathrm{X} \times 2+1\) of VSR[ XT] are undefined.
3. \(\mathrm{FR}, \mathrm{FI}\), and FPRF are not modified.

For the VSX Vector Convert Single-Precision to Signed Integer Doubleword (xvcvspsxd) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1 .
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXCVI & (if Invalid Integer Convert)
\end{tabular}
2. OXTFFFFFFF_FFF_FFFF is placed into doubleword element i of VSR[XT] if the single-precision operand in word element ix2 of VSR[XB] is a positive number or +Infinity.
\(0 \times 8000 \_0000 \_0000 \_0000\) is placed into doubleword element \(i\) of VSR[XT] if the single-precision operand in word element \(\mathrm{i} \times 2\) of \(\mathrm{VSR}[\mathrm{XB}]\) is a negative number, - Infinity, or NaN.
3. \(F R, F I\), and \(F P R F\) are not modified.

For the VSX Vector Convert Single-Precision to Unsigned Integer Doubleword (xvcvspuxd) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1.

VXSNAN (if SNaN )

VXCVI (if Invalid Integer Convert)
2. OXFFFF_FFF_ FFFF_FFF \(^{2}\) is placed into doubleword element \(i\) of VSR[XT] if the single-precision operand in word element \(\bar{x} \times 2\) of \(\operatorname{VSR}[X B]\) is a positive number or +Infinity.
\(0 \times 0000 \_0000 \_0000 \_0000\) is placed into doubleword element \(i\) of VSR[ XT] if the single-precision operand in word element \(\bar{i}^{-} \times 2\) of VSR[XB] is a negative number, - Infinity, or NaN .
3. \(F R, F I\), and \(F P R F\) are not modified.

For the VSX Vector Convert Single-Precision to Signed Integer Word (xvcvspsxw) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1.
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXCVI & (if Invalid Integer Convert)
\end{tabular}
2. OXTFFF_FFFF is placed into word element \(i\) of VSR[XT] if the single-precision operand in word element \(i\) of VSR[XB] is a positive number or +Infinity.
\(0 \times 8000 \_0000\) is placed into word element \(i\) of VSR[ XT] if the single-precision operand in word element \(i\) of VSR [ \(X B]\) is a negative number, - Infinity, or NaN.

The contents of word element \(2 \times i+1\) of VSR[ XT] are undefined.
3. \(F R, F I\), and \(F P R F\) are not modified.

For the VSX Vector Convert Single-Precision to Unsigned Integer Word (xvcvspuxw) instruction:
1. One or two of the following Invalid Operation exceptions are set to 1.
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXCVI & (if Invalid Integer Convert)
\end{tabular}
2. OXFFF_FFFF is placed into word element \(i\) of VSR[XT] if the single-precision operand in the corresponding word element 2 xi of VSR[ XB\(]\) is a positive number or +Infinity.
\(0 \times 00000000\) is placed into word element i of VSR[ XT] if the single-precision operand in word element 2 xi of VSR[ XB] is a negative number, - Infinity, or NaN.

The contents of word element \(2 x i+1\) of VSR[ XT] are undefined.
3. \(F R, F I\), and \(F P R F\) are not modified.

For the VSX Scalar Floating-Point Compare instructions:
1. One or two of the following Invalid Operation exceptions are set to 1.
\begin{tabular}{ll} 
VXSNAN & (if SNaN) \\
VXCVI & (if Invalid Integer Convert)
\end{tabular}
2. \(F R, F I\) and \(C\) are unchanged.
3. FPCC is set to reflect unordered.

For the VSX Vector Compare Single-Precision instructions:
1. One or two of the following Invalid Operation exceptions are set to 1 .

VXSNAN (if SNaN )

\section*{Version 3.0}

VXCVI (if Invalid Integer Convert)
2. \(0 \times 0000 \_0000\) is placed into its respective word element of VSR[XT].
3. \(F R, F I\), and \(F P R F\) are not modified.

For the vector double-precision compare instructions:
1. One or two of the following Invalid Operation exceptions are set to 1 .
\begin{tabular}{ll} 
VXSNAN & (if SNaN ) \\
VXCVI & (if Invalid Integer Convert)
\end{tabular}
2. \(0 \times 00000_{-} 0000_{-} 00000_{-} 0000\) is placed into its respective doubleword element of VSR[ XT].
3. \(F R, F I\), and \(F P R F\) are not modified.

For VSX Vector Convert Single-Precision to Half-Precision with round (xscvsphp), do the following.
1. VXSNAN is set to 1 .
2. The half-precision representation of a Quiet NaN is placed into the rightmost halfword of its respective word element of VSR[XT]. The contents of the leftmost halfword of its respective word element of VSR[XT] are set to 0 .
3. \(F R, F I\), and \(F P R F\) are not modified.

For VSX Vector Convert Half-Precision to Single-Precision (xscvhpsp), do the following.
1. VXSNAN is set to 1 .
2. The half-precision representation of a Quiet NaN is placed into the rightmost halfword of its respective word element of VSR[XT]. The contents of the leftmost halfword of its respective word element of VSR[ \(X T]\) are set to 0.
3. \(F R, F I\), and \(F P R F\) are not modified.

\subsection*{7.4.2 Floating-Point Zero Divide Exception}

\subsection*{7.4.2.1 Definition}

A Zero Divide exception occurs when a VSX Floating-Point Divide \({ }^{[1]}\) instruction is executed with a zero divisor value and a finite nonzero dividend value.

A Zero Divide exception also occurs when a VSX Floating-Point Reciprocal Estimate \({ }^{[2]}\) instruction or a VSX Floating-Point Reciprocal Square Root Estimate \({ }^{[3]}\) instruction is executed with an operand value of zero.

The action to be taken depends on the setting of the Zero Divide Exception Enable bit of the FPSCR.

\subsection*{7.4.2.2 Action for ZE=1}

When Zero Divide exception is enabled \((Z E=1)\) and a Zero Divide exception occurs, the following actions are taken:
For any of the following instructions,
VSX Scalar Floating-Point Divide instructions: xsdivdp, xsdivsp

VSX Scalar Floating-Point Reciprocal Estimate instructions xsredp, xsresp

VSX Scalar Floating-Point Reciprocal Square Root Estimate instructions xsrsqrtedp, xsrsqrtesp
do the following.
1. \(Z X\) is set to 1 .
2. Update of VSR[ XT] is suppressed.
3. FR and Fl are set to 0 .
4. \(F P R F\) is unchanged.

For VSX Scalar Divide Quad-Precision (xsdivqp), do the following.
1. \(Z X\) is set to 1 .
2. Update of VSR[VRT+32] is suppressed.
3. \(F R\) and \(F I\) are set to \(0 . F P R F\) is not modified.

For any of the following instructions,
VSX Vector Floating-Point Divide instructions
xsdivdp, xsdivsp, xvdivdp, xvdivsp
VSX Vector Floating-Point Reciprocal Estimate instructions
xsredp, xsresp, xvredp, xvresp
VSX Vector Floating-Point Reciprocal Square Root Estimate instructions xsrsqrtedp, xsrsqrtesp, xvrsqrtedp, xvrsqrtesp

\footnotetext{
1. VSX Vector Floating-Point Divide instructions: xsdivdp, xsdivsp, xvdivdp, xvdivsp
2. VSX Floating-Point Reciprocal Estimate instructions: xsredp, xsresp, xvredp, xvresp
3. VSX Floating-Point Reciprocal Square Root Estimate instructions: xsrsqrtedp, xsrsqrtesp, xvrsqrtedp, xvrsqrtesp
}
do the following.
1. \(Z X\) is set to 1 .
2. Update of VSR[ XT] is suppressed for all vector elements.
3. \(F R\) and \(F I\) are unchanged.
4. \(F P R F\) is unchanged.

\subsection*{7.4.2.3 Action for \(Z E=0\)}

When Zero Divide exception is disabled \((Z E=0)\) and a Zero Divide exception occurs, the following actions are taken:
For VSX Scalar Floating-Point Divide \({ }^{[1]}\) instructions, do the following.
1. \(Z X\) is set to 1 .
2. An Infinity, having a sign determined by the XOR of the signs of the source operands, is placed into doubleword element 0 of VSR[ XT] in double-precision format. The contents of doubleword element 1 of VSR[ \(X T]\) are undefined.
3. \(F R\) and \(F I\) are set to 0 .
4. FPRF is set to indicate the class and sign of the result ( \(\pm\) Infinity).

For VSX Scalar Divide Quad-Precision (xsdivqp), do the following.
1. \(Z X\) is set to 1 .
2. An Infinity, having a sign determined by the XOR of the signs of the source operands, is placed into VSR[ VRT +32 ] in quad-precision format.
3. \(F R\) and \(F I\) are set to \(0 . F P R F\) is set to indicate the class and sign of the result ( \(\pm\) Infinity).

For VSX Vector Divide Double-Precision (xvdivdp), do the following.
1. \(Z X\) is set to 1 .
2. For each vector element causing a Zero Divide exception, an Infinity, having a sign determined by the XOR of the signs of the source operands, is placed into its respective doubleword element of VSR[ XT] in double-precision format.
3. \(F R, F I\), and \(F P R F\) are not modified.

For VSX Vector Divide Single-Precision (xvdivsp), do the following.
1. \(Z X\) is set to 1 .
2. For each vector element causing a Zero Divide exception, an Infinity, having a sign determined by the XOR of the signs of the source operands, is placed into its respective word element of VSR[XT] in single-precision format.
3. \(F R, F I\), and \(F P R F\) are not modified.

\footnotetext{
1. VSX Scalar Floating-Point Divide instructions: xsdivdp, xsdivsp
}

For VSX Scalar Floating-Point Reciprocal Estimate \({ }^{[1]}\) instructions and VSX Scalar Floating-Point Reciprocal Square Root Estimate \({ }^{[2]}\) instructions, do the following.
1. \(Z X\) is set to 1 .
2. An Infinity, having the sign of the source operand, is placed into doubleword element 0 of VSR[ XT] in double-precision format. The contents of doubleword element 1 of VSR[ XT] are undefined.
3. FR and Fl are set to 0 .
4. \(\operatorname{FPRF}\) is set to indicate the class and sign of the result ( \(\pm\) Infinity).

For the VSX Vector Reciprocal Estimate Double-Precision (xvredp) and VSX Vector Reciprocal Square Root Estimate Double-Precision (xvrsqrtedp) instructions:
1. \(Z X\) is set to 1 .
2. For each vector element causing a Zero Divide exception, an Infinity, having the sign of the source operand, is placed into its respective doubleword element of VSR[ XT] in double-precision format.
3. \(F R, F I\), and \(F P R F\) are not modified.

For the VSX Vector Reciprocal Estimate Single-Precision (xvresp) and VSX Vector Reciprocal Square Root Estimate Single-Precision (xvrsqrtesp) instructions:
1. \(Z X\) is set to 1 .
2. For each vector element causing a Zero Divide exception, an Infinity, having the sign of the source operand, is placed into its respective word element of VSR[ XT] in single-precision format.
3. \(F R, F I\), and \(F P R F\) are not modified.

\footnotetext{
1. VSX Scalar Floating-Point Reciprocal Estimate instructions: xsredp, xsresp
2. VSX Scalar Floating-Point Reciprocal Square Root Estimate instructions: xsrsqrtedp, xsrsqrtesp
}

\subsection*{7.4.3 Floating-Point Overflow Exception}

\subsection*{7.4.3.1 Definition}

An Overflow exception occurs when the magnitude of what would have been the rounded result if the exponent range were unbounded exceeds that of the largest finite number of the specified result precision.

The action to be taken depends on the setting of the Overflow Exception Enable bit of the FPSCR.

\subsection*{7.4.3.2 Action for OE=1}

When Overflow exception is enabled (OE=1) and an Overflow exception occurs, the following actions are taken:
For the VSX Vector round and Convert Double-Precision to Single-Precision format (xscvdpsp) instruction:
1. OX is set to 1 .
2. If the unbiased exponent of the normalized intermediate result is less than or equal to 318 (Emax+192), the exponent is adjusted by subtracting 192. Otherwise the result is undefined.
3. The adjusted rounded result is placed into word element 0 of \(\operatorname{VSR}[X T]\) in single-precision format. The contents of word elements 1-3 of VSR[XT] are undefined.
4. Unless the result is undefined, FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

For VSX Scalar Double-Precision Arithmetic \({ }^{[1]}\) instructions, do the following.
1. OX is set to 1 .
2. The exponent of the normalized intermediate result is adjusted by subtracting 1536.
3. The adjusted rounded result is placed into doubleword element 0 of \(\operatorname{VSR}[X T]\) in double-precision format. The contents of doubleword element 1 of VSR[XT] are undefined.
4. FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

For VSX Scalar Single-Precision Arithmetic \({ }^{[2]}\) instructions, do the following.
1. OX is set to 1 .
2. The exponent is adjusted by subtracting 192.
3. The adjusted and rounded result is placed into doubleword element 0 of VSR[XT] in double-precision format. The contents of doubleword element 1 of VSR[XT] are undefined.
4. FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

\footnotetext{
1. VSX Scalar Double-Precision Arithmetic instructions:
xsadddp, xsdivdp, xsmuldp, xsredp, xssubdp, xsmaddadp, xsmaddmdp, xsmsubadp, xsmsubmdp, xsnmaddadp, xsnmaddmdp, xsnmsubadp, xsnmsubmdp
2. VSX Scalar Single-Precision Arithmetic instructions:
xsaddsp, xsdivsp, xsmulsp, xsresp, xssubsp, xsmaddasp, xsmaddmsp, xsmsubasp, xsmsubmsp, xsnmaddasp, xsnmaddmsp, xsnmsubasp, xsnmsubmsp
}

For any of the following instruction classes,
VSX Scalar Quad-Precision Arithmetic instructions:
xsaddqp[o], xsdivqp[o], xsmulqp[o], xssqrtqp[o], xssubqp[o]
xsmaddqp[o], xsmsubqp[o], xsnmaddqp[o], xsnmsubqp[o]
VSX Scalar Round Quad-Precision to Double-Extended-Precision (xsrqpxp)
do the following.
1. \(O X\) is set to 1 .
2. The exponent is adjusted by subtracting 24576.
3. The adjusted, rounded result is placed into VSR[ VRT +32 ] in quad-precision format.
4. Unless the result is undefined, FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

For VSX Scalar Convert Quad-Precision to Double-Precision [using round to Odd] (xscvqpdp), do the following.
1. OX is set to 1 .
2. The exponent is adjusted by subtracting 1536. If the adjusted exponent is greater than +1023 ( Emax ), the result is undefined.
3. The adjusted, rounded result is placed into doubleword element 0 of VSR[VRT+32] in double-precision format.
\(0 \times 00000_{-0000} 0000 \_0000\) is placed into doubleword element 1 of VSR[VRT+32].
4. Unless the result is undefined, FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

For VSX Scalar Convert Double-Precision to Half-Precision with round (xscvdphp), do the following.
1. \(O X\) is set to 1 .
2. The exponent is adjusted by subtracting 24. If the adjusted exponent is greater than +15 ( Emax ), the result is undefined.
3. The adjusted, rounded result is placed into rightmost halfword of doubleword element 0 of VSR[ XT] in half-precision format.

The contents of the leftmost 3 halfwords of doubleword element 0 of VSR[ XT] are set to 0 .
The contents of doubleword element 1 of VSR[ XT] are undefined.
4. Unless the result is undefined, FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

For VSX Vector Double-Precision Arithmetic \({ }^{[1]}\) instructions, VSX Vector Single-Precision Arithmetic \({ }^{[2]}\) instructions, and VSX Vector round and Convert Double-Precision to Single-Precision format instruction (xvcvdpsp), do the following.
1. OX is set to 1 .
2. Update of VSR[XT] is suppressed for all vector elements.
3. FR, FI, and FPRF are not modified.

For VSX Vector Convert Single-Precision to Half-Precision with round (xvcvsphp), do the following.
1. \(O X\) is set to 1 .
2. VSR[ \(X T]\) is not modified.
3. \(F R, F I\), and FPRF are not modified.

\footnotetext{
1. VSX Vector Double-Precision Arithmetic instructions:
xvadddp, xvdivdp, xvmuldp, xvredp, xvsubdp, xvmaddadp, xsmaddmdp, xvmsubadp, xvmsubmdp, xvnmaddadp, xvnmaddmdp, xvnmsubadp, xvnmsubmdp
2. VSX Vector Single-Precision Arithmetic instructions:
xvaddsp, xvdivsp, xvmulsp, xvresp, xvsubsp, xvmaddasp, xvmaddmsp, xvsmsubasp, xvmsubmsp, xvnmaddasp, xvnmaddmsp, xvnmsubasp, xvnmsubmsp
}

\subsection*{7.4.3.3 Action for \(\mathrm{OE}=0\)}

When Overflow exception is disabled \((0 E=0)\) and an Overflow exception occurs, the following actions are taken:
1. \(O X\) and \(X X\) are set to 1 .
2. The result is determined by the rounding mode ( RN ) and the sign of the intermediate result as follows:

\section*{Round to Nearest Even}

For negative overflow, the result is - Infinity.
For positive overflow, the result is +Infinity.

\section*{Round toward Zero}

For negative overflow, the result is the format's most negative finite number.
For positive overflow, the result is the format's most positive finite number.

\section*{Round toward +Infinity}

For negative overflow, the result is the format's most negative finite number.
For positive overflow, the result is +Infinity.

\section*{Round toward - Infinity}

For negative overflow, the result is -Infinity.
For positive overflow, the result is the format's most positive finite number.
For VSX Scalar round and Convert Double-Precision to Single-Precision format (xscvdpsp):
3. The result is placed into word element 0 of VSR[XT] as a single-precision value. The contents of word elements 1-3 of VSR[ XT] are undefined.
4. \(F R\) is undefined.
5. Fl is set to 1 .
6. \(\operatorname{FPRF}\) is set to indicate the class and sign of the result.

For VSX Scalar Double-Precision Arithmetic \({ }^{[1]}\) instructions and VSX Scalar Single-Precision Arithmetic \({ }^{[2]}\) instructions, do the following.
3. The result is placed into doubleword element 0 of \(\operatorname{VSR}[X T]\) as a double-precision value. The contents of doubleword element 1 of \(\operatorname{VSR}[X T]\) are undefined.
4. \(F R\) is undefined.
5. Fl is set to 1 .
6. FPRF is set to indicate the class and sign of the result.

For any of the following instructions,
VSX Scalar Quad-Precision Arithmetic instructions:
xsaddqp[o], xsdivqp[o], xsmulqp[o], xssubqp[o]
xsmaddqp[o], xsmsubqp[o], xsnmaddqp[o], xsnmsubqp[o]
VSX Scalar Quad-Precision Round to Double-Extended-Precision (xsrqpxp)

\footnotetext{
1. VSX Scalar Double-Precision Arithmetic instructions: xsadddp, xsdivdp, xsmuldp, xsredp, xssubdp, xsmaddadp, xsmaddmdp, xsmsubadp, xsmsubmdp, xsnmaddadp, xsnmaddmdp, xsnmsubadp, xsnmsubmdp
2. VSX Scalar Single-Precision Arithmetic instructions:
xsaddsp, xsdivsp, xsmulsp, xsresp, xssubsp, xsmaddasp, xsmaddmsp, xsmsubasp, xsmsubmsp, xsnmaddasp, xsnmaddmsp, xsnmsubasp, xsnmsubmsp
}
do the following.
3. The result is placed into VSR[ VRT +32 ] in quad-precision format.
4. \(F R\) is undefined. \(F I\) is set to \(1 . F P R F\) is set to indicate the class and sign of the result.

For VSX Scalar Convert Quad-Precision to Double-Precision (xscvqpdp), do the following.
3. The result is placed into doubleword element 0 of VSR[VRT+32] as a double-precision value. \(0 \times 00000^{0} 0000_{-} 0000 \_0000\) is placed into doubleword element 1 of VSR[VRT+32].
4. \(F R\) is undefined. \(F I\) is set to 1. FPRF is set to indicate the class and sign of the result.

For VSX Scalar Convert Double-Precision to Half-Precision (xscvdphp), do the following.
1. \(O X\) and \(X X\) are set to 1 .
2. The result is placed into the rightmost halfword of doubleword element 0 of VSR[ XT] as a half-precision value.

The contents of the leftmost 3 halfwords of doubleword element 0 of VSR[ XT] are set to 0 .
The contents of doubleword element 1 of VSR[ XT] are undefined.
3. \(F R\) is undefined. \(F I\) is set to 1. FPRF is set to indicate the class and sign of the result.

For VSX Vector Double-Precision Arithmetic \({ }^{[1]}\) instructions, do the following.
3. For each vector element causing an Overflow exception, the result is placed into its respective doubleword element of VSR[XT] in double-precision format.
4. FR, FI, and FPRF are not modified.

For VSX Vector Single-Precision Arithmetic \({ }^{[2]}\) instructions and VSX Vector round and Convert Double-Precision to Single-Precision format (xvcvdpsp), do the following.
3. For each vector element causing an Overflow exception, the result is placed into its respective word element of VSR[XT] in single-precision format.
4. FR, FI, and FPRF are not modified.

For VSX Vector Convert Single-Precision to Half-Precision with round (xvcvsphp), do the following.
1. \(O X\) and \(X X\) are set to 1 .
2. For each vector element causing an Overflow exception, the result is placed into the rightmost halfword of its respective word element of VSR[ XT] in half-precision format.

The contents of the leftmost halfword of its respective word element of VS R[XT] are set to 0 .
3. \(F R, F I\), and \(F P R F\) are not modified.

\footnotetext{
1. VSX Vector Double-Precision Arithmetic instructions:
xvadddp, xvdivdp, xvmuldp, xvredp, xvsubdp, xvmaddadp, xvmaddmdp, xvmsubadp, xvmsubmdp, xvnmaddadp, xvnmaddmdp, xvnmsubadp, xvnmsubmdp
2. VSX Vector Single-Precision Arithmetic instructions:
xvaddsp, xvdivsp, xvmulsp, xvresp, xvsubsp, xvmaddasp, xvmaddmsp, xvmsubasp, xvmsubmsp, xvnmaddasp, xvnmaddmsp, xvnmsubasp, xvnmsubmsp
}

\section*{410} Power ISA \({ }^{\text {TM }}\) - Book I

\subsection*{7.4.4 Floating-Point Underflow Exception}

\subsection*{7.4.4.1 Definition}

Underflow exception is defined separately for the enabled and disabled states:

\section*{Enabled:}

Underflow occurs when the intermediate result is "Tiny".

\section*{Disabled:}

Underflow occurs when the intermediate result is "Tiny" and there is "Loss of Accuracy".
A tiny result is detected before rounding, when a nonzero intermediate result computed as though both the precision and the exponent range were unbounded would be less in magnitude than the smallest normalized number.

If the intermediate result is tiny and Underflow exception is disabled \((U E=0)\), the intermediate result is denormalized (see Section 7.3.2.4 , "Normalization and Denormalization" on page 379) and rounded (see Section 7.3.2.6 , "Rounding" on page 383) before being placed into the target VSR.

Loss of accuracy is detected when the delivered result value differs from what would have been computed were both the precision and the exponent range unbounded.

The action to be taken depends on the setting of the Underflow Exception Enable bit of the FPSCR.

\subsection*{7.4.4.2 Action for UE=1}

When Underflow exception is enabled \((U E=1)\) and an Underflow exception occurs, the following actions are taken:
For VSX Scalar round and Convert Double-Precision to Single-Precision format (xscvdpsp), do the following.
1. \(U X\) is set to 1 .
2. If the unbiased exponent of the normalized intermediate result is greater than or equal to -319 ( Emi n -192), the exponent is adjusted by adding 192. Otherwise the result is undefined.
3. The adjusted rounded result is placed into word element 0 of VSR[ XT] in single-precision format. The contents of word elements 1-3 of VSR[ XT] are undefined.
4. Unless the result is undefined, FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

For VSX Scalar Double-Precision Arithmetic \({ }^{[1]}\) instructions and VSX Scalar Double-Precision Reciprocal Estimate (xsredp), do the following.
1. UX is set to 1 .
2. The exponent of the normalized intermediate result is adjusted by adding 1536.
3. The adjusted rounded result is placed into doubleword element 0 of VSR[XT] in double-precision format. The contents of doubleword element 1 of VSR[ XT] are undefined.
4. \(\quad \mathrm{FPRF}\) is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

\footnotetext{
1. VSX Scalar Double-Precision Arithmetic instructions:
xsadddp, xsdivdp, xsmuldp, xssubdp, xsmaddadp, xsmaddmdp, xsmsubadp, xsmsubmdp, xsnmaddadp, xsnmaddmdp, xsnmsubadp, xsnmsubmdp
}

For any of the following instructions,
VSX Scalar Quad-Precision Arithmetic instructions:
xsaddqp[o], xsdivqp[o], xsmulqp[o], xssubqp[o]
xsmaddqp[o], xsmsubqp[o], xsnmaddqp[o], xsnmsubqp[o]
VSX Scalar Round Quad-Precision to Double-Extended-Precision (xsrqpxp)
do the following.
1. \(U X\) is set to 1 .
2. The exponent of the normalized intermediate result is adjusted by adding 24576 .
3. The adjusted, rounded result is placed into VSR[ VRT +32 ] in quad-precision format.
4. Unless the result is undefined, FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

For VSX Scalar Convert Quad-Precision to Double-Precision [using round to Odd] (xscvqpdp[o]), do the following.
1. \(U X\) is set to 1 .
2. The exponent of the normalized intermediate result is adjusted by adding 1536. If the adjusted exponent is less than - 1022 , the result is undefined.
3. The adjusted, rounded result is placed into doubleword element 0 of VSR[ VRT +32 ] in double-precision format.
\(0 \times 0000 \_0000 \_0000 \_0000\) is placed into doubleword element 1 of VSR[VRT +32 ].
4. Unless the result is undefined, FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

For VSX Scalar Single-Precision Arithmetic \({ }^{[1]}\) instructions and VSX Scalar Single-Precision Reciprocal Estimate (xsresp), do the following.
1. UX is set to 1 .
2. The exponent is adjusted by adding 192.
3. The adjusted rounded result is placed into doubleword element 0 of \(\operatorname{VSR}[X T]\) in double-precision format. The contents of doubleword element 1 of VSR[XT] are undefined.
4. FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

\section*{Programming Note}

The FR and FI bits are provided to allow the system floating-point enabled exception error handler, when invoked because of an Underflow exception, to simulate a "trap disabled" environment. That is, the FR and FI bits allow the system floating-point enabled exception error handler to unround the result, thus allowing the result to be denormalized and correctly rounded.

For VSX Scalar Convert Double-Precision to Half-Precision with round (xscvdphp), do the following.
1. \(U X\) is set to 1 .

\footnotetext{
1. VSX Scalar Single-Precision Arithmetic instructions:
}
xsaddsp, xsdivsp, xsmulsp, xssubsp, xsmaddasp, xsmaddmsp, xsmsubasp, xsmsubmsp, xsnmaddasp, xsnmaddmsp, xsnmsubasp, xsnmsubmsp

\section*{412 Power ISA \({ }^{\text {TM }}\) - Book I}
2. The exponent of the normalized intermediate result is adjusted by adding 24. If the adjusted exponent is less than - 14 , the result is undefined.
3. The adjusted, rounded result is placed into rightmost halfword of doubleword element 0 of VSR[ XT] in half-precision format.

The contents of the leftmost 3 halfwords of doubleword element 0 of VSR[ XT] are set to 0 .
The contents of doubleword element 1 of VSR[ XT] are undefined.
4. Unless the result is undefined, FPRF is set to indicate the class and sign of the result ( \(\pm\) Normal Number).

For VSX Vector Floating-Point Arithmetic \({ }^{[1]}\) instructions, VSX Vector Floating-Point Reciprocal Estimate \({ }^{[2]}\) instructions, and VSX Vector round and Convert Double-Precision to Single-Precision format (xvcvdpsp), do the following.
1. UX is set to 1 .
2. Update of \(\mathrm{VSR}[\mathrm{XT}]\) is suppressed for all vector elements.
3. FR, FI, and FPRF are not modified.

For VSX Vector Convert Single-Precision to Half-Precision with round (xvcvsphp), do the following.
1. \(U X\) is set to 1 .
2. VSR[ \(X T]\) is not modified.
3. \(F R, F I\), and \(F P R F\) are not modified.

\subsection*{7.4.4.3 Action for UE=0}

When Underflow exception is disabled (UE=0) and an Underflow exception occurs, the following actions are taken:
For VSX Scalar round and Convert Double-Precision to Single-Precision format (xscvdpsp), do the following.
1. UX is set to 1 .
2. The result is placed into word element 0 of \(\operatorname{VSR}[\mathrm{XT}]\) in single-precision format. The contents of word elements 1-3 of VSR[XT] are undefined.
3. FPRF is set to indicate the class and sign of the result.

For VSX Scalar Floating-Point Arithmetic \({ }^{[3]}\) instructions and VSX Scalar Reciprocal Estimate \({ }^{[4]}\) instructions, do the following.
1. UX is set to 1 .
2. The result is placed into doubleword element 0 of \(\operatorname{VSR}[X T]\) in double-precision format. The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

\footnotetext{
1. VSX Vector Arithmetic instructions: xvadddp, xvdivdp, xvmuldp, xvsubdp, xvaddsp, xvdivsp, xvmulsp, xvsubsp, xvmaddadp, xvmaddmdp, xvmsubadp, xvmsubmdp, xvnmaddadp, xvnmaddmdp, xvnmsubadp, xvnmsubmdp, xvmaddasp, xvmaddmsp, xvmsubasp, xvmsubmsp, xvnmaddasp, xvnmaddmsp, xvnmsubasp, xvnmsubmsp
2. VSX Vector Floating-Point Reciprocal Estimate instructions: xvredp, xvresp
3. VSX Scalar Floating-Point Arithmetic instructions: xsadddp, xsdivdp, xsmuldp, xssubdp, xsaddsp, xsdivsp, xsmulsp, xssubsp, xsmaddadp, xsmaddmdp, xsmsubadp, xsmsubmdp, xsnmaddadp, xsnmaddmdp, xsnmsubadp, xsnmsubmdp, xsmaddasp, xsmaddmsp, xsmsubasp, xsmsubmsp, xsnmaddasp, xsnmaddmsp, xsnmsubasp, xsnmsubmsp
4. VSX Scalar Reciprocal Estimate instructions: xsredp, xsresp
}
3. FPRF is set to indicate the class and sign of the result.

For any of the following instructions,
VSX Scalar Quad-Precision Arithmetic instructions:
xsaddqp[o], xsdivqp[o], xsmulqp[o], xssubqp[o]
xsmaddqp[o], xsmsubqp[o], xsnmaddqp[o], xsnmsubqp[o]
VSX Scalar Round Quad-Precision to Double-Extended-Precision (xsrqpxp)
do the following.
1. \(U X\) is set to 1 .
2. The result is placed into VSR[ VRT +32 ] in quad-precision format.
3. \(F\) PRF is set to indicate the class and sign of the result.

For VSX Scalar Convert Quad-Precision to Double-Precision (xscvqpdp), do the following.
1. \(U X\) is set to 1 .
2. The result is placed into doubleword element 0 of VSR[ VRT+32] in double-precision format. \(0 \times 0000_{-} 0000 \_0000 \_0000\) is placed into doubleword element 1 of VSR[VRT+32].
3. \(F P R F\) is set to indicate the class and sign of the result.

For VSX Scalar Convert Double-Precision to Half-Precision with round (xscvdphp), do the following.
1. \(U X\) is set to 1 .
2. The result is placed into the rightmost halfword of doubleword element 0 of \(\operatorname{VSR}[X T]\) as a half-precision value.

The contents of the leftmost 3 halfwords of doubleword element 0 of VSR[ XT] are set to 0 .
The contents of doubleword element 1 of VSR[ XT] are undefined.
3. \(\operatorname{FPRF}\) is set to indicate the class and sign of the result.

For VSX Vector Double-Precision Arithmetic \({ }^{[1]}\) instructions and VSX Vector Reciprocal Estimate Double-Precision (xvredp), do the following.
1. UX is set to 1 .
2. For each vector element causing an Underflow exception, the result is placed into its respective doubleword element of VSR[XT] in double-precision format.
3. FR, FI, and FPRF are not modified.

For VSX Vector Single-Precision Arithmetic \({ }^{[2]}\) instructions, VSX Vector Reciprocal Estimate Single-Precision (xvresp), and VSX Vector round and Convert Double-Precision to Single-Precision format (xvcvdpsp), do the following.
1. UX is set to 1 .
1. VSX Vector Double-Precision Arithmetic instructions:
xvadddp, xvdivdp, xvmuldp, xvsubdp, xvmaddadp, xvmaddmdp, xvmsubadp, xvmsubmdp, xvnmaddadp, xvnmaddmdp, xvnmsubadp, xvnmsubmdp
2. VSX Vector Single-Precision Arithmetic instructions:
xvaddsp, xvdivsp, xvmulsp, xvsubsp, xvmaddasp, xvmaddmsp, xvmsubasp, xvmsubmsp, xvnmaddasp, xvnmaddmsp, xvnmsubasp, xvnmsubmsp
2. For each vector element causing an Underflow exception, the result is placed into its respective word element of VSR[XT] in single-precision format.
3. FR, FI, and FPRF are not modified.

For VSX Vector Convert Single-Precision to Half-Precision with round (xvcvsphp), do the following.
1. UX is set to 1 .
2. For each vector element causing an Underflow exception, the result is placed into the rightmost halfword of its respective word element of VSR[ XT] in half-precision format.

The contents of the leftmost halfword of its respective word element of VSR[ XT] are set to 0 .
3. \(F R, F I\), and \(F P R F\) are not modified.

\subsection*{7.4.5 Floating-Point Inexact Exception}

\subsection*{7.4.5.1 Definition}

An Inexact exception occurs when one of two conditions occur during rounding:
1. The rounded result differs from the intermediate result assuming both the precision and the exponent range of the intermediate result to be unbounded. In this case the result is said to be inexact. (If the rounding causes an enabled Overflow exception or an enabled Underflow exception, an Inexact exception also occurs only if the significands of the rounded result and the intermediate result differ.)
2. The rounded result overflows and Overflow exception is disabled.

The action to be taken depends on the setting of the Inexact Exception Enable bit of the FPSCR.

\subsection*{7.4.5.2 Action for \(X E=1\)}

\section*{Programming Note}

In some implementations, enabling Inexact exceptions can degrade performance more than does enabling other types of floating-point exception.

When Inexact exception is enabled (UE=1) and an Inexact exception occurs, the following actions are taken:
For the VSX Vector round and Convert Double-Precision to Single-Precision format (xscvdpsp) instruction:
1. \(X X\) is set to 1 .
2. The result is placed into word element 0 of \(\operatorname{VSR}[X T]\) in single-precision format. The contents of word elements 1-3 of VSR[XT] are undefined.
3. FPRF is set to indicate the class and sign of the result.

For VSX Scalar Floating-Point Arithmetic \({ }^{[1]}\) instructions, VSX Scalar Round to Double-Precision Integer Exact using Current rounding mode (xsrdpic), and VSX Scalar Integer to Floating-Point Format Conversion \({ }^{[2]}\) instructions, do the following.
1. XX is set to 1 .
2. The result is placed into doubleword element 0 of \(\mathrm{VSR}[X T]\) in double-precision format. The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.
3. FPRF is set to indicate the class and sign of the result.

For VSX Scalar Floating-Point to Integer Word Format Conversion \({ }^{[3]}\) instructions, do the following.
1. XX is set to 1 .
2. The result is placed into word element 1 of VSR[XT]. The contents of word elements 0 , 2, and 3 of VSR[XT] are undefined.
3. FPRF is set to indicate the class and sign of the result.

\footnotetext{
1. VSX Scalar Floating-Point Arithmetic instructions:
xsadddp, xsdivdp, xsmuldp, xssubdp, xsaddsp, xsdivsp, xsmulsp, xssubsp, xsmaddadp, xsmaddmdp, xsmsubadp, xsmsubmdp, xsnmaddadp, xsnmaddmdp, xsnmsubadp, xsnmsubmdp, xsmaddasp, xsmaddmsp, xsmsubasp, xsmsubmsp, xsnmaddasp, xsnmaddmsp, xsnmsubasp, xsnmsubmsp
2. VSX Scalar Integer to Floating-Point Format Conversion instructions: xscvsxddp, xscvuxddp, xscvsxdsp, xscvuxdsp
3. VSX Scalar Floating-Point to Integer Word Format Conversion instructions: xscvdpsxws, xscvdpuxws
}

For any of the following instructions,
VSX Scalar Quad-Precision Arithmetic instructions:
xsaddqp[o], xsdivqp[o], xsmulqp[o], xssqrtqp[o], xssubqp[o]
xsmaddqp[o], xsmsubqp[o], xsnmaddqp[o], xsnmsubqp[o]
VSX Scalar Quad-Precision Round instructions:
xsrqpi, xsrqpxp
do the following.
1. \(X X\) is set to 1 .
2. The result is placed into VSR[VRT +32 ] in quad-precision format.
3. \(F R\) is set to indicate if the rounded result was incremented. Fl is set to 1 . FPRF is set to indicate the class and sign of the result.

For VSX Scalar Convert Quad-Precision to Double-Precision (xscvqpdp), do the following.
1. \(X X\) is set to 1 .
2. The result is placed into doubleword element 0 of VSR[ VRT +32 ] in double-precision format.
\(0 \times 0000^{2} 00000^{2} 0000_{-} 0000\) is placed into doubleword element 1 of VSR[VRT+32].
3. \(F R\) is set to indicate if the rounded result was incremented. \(F I\) is set to 1 . \(F P R F\) is set to indicate the class and sign of the result.

For VSX Scalar truncate \& Convert Quad-Precision to Signed Doubleword (xscvqpsdz), do the following.
1. \(X X\) is set to 1 .
2. The result is placed into doubleword element 0 of VSR[ \(X T]\) in signed integer format. \(0 \times 0000^{2} 00000^{2} 0000_{-} 0000\) is placed into doubleword element 1 of VSR[VRT+32].
3. \(F R\) is set to \(0 . F I\) is set to \(1 . F P R F\) is undefined.

For VSX Scalar truncate \& Convert Quad-Precision to Signed Word (xscvqpswz), do the following.
1. \(X X\) is set to 1 .
2. The result is placed into word element 1 of VSR[ XT] in signed integer format. \(0 \times 0000 \_0000\) is placed into word elements 0,2 , and 3 of VSR[ VRT +32 ].
3. \(F R\) is set to \(0 . F I\) is set to 1 . \(F P R F\) is undefined.

For VSX Scalar truncate \& Convert Quad-Precision to Unsigned Doubleword (xscvqpudz), do the following.
1. \(X X\) is set to 1 .
2. The result is placed into doubleword element 0 of \(V S R[X T]\) in unsigned integer format. \(0 \times 0000 \_0000 \_0000 \_0000\) is placed into doubleword element 1 of VSR[ VRT+32].
3. \(F R\) is set to \(0 . F I\) is set to 1 . \(F P R F\) is undefined.

For VSX Scalar truncate \& Convert Quad-Precision to Unsigned Word (xscvqpuwz), do the following.
1. \(X X\) is set to 1 .
2. The result is placed into word element 1 of VSR[ \(X T]\) in unsigned integer format.
\(0 \times 0000 \_0000\) is placed into word elements 0,2 , and 3 of VSR[ VRT +32 ].
3. \(F R\) is set to \(0 . F \mid\) is set to \(1 . F P R F\) is undefined.

For VSX Scalar Convert Double-Precision to Half-Precision with round (xscvdphp), do the following.
1. \(X X\) is set to 1 .
2. The result is placed into the rightmost halfword of doubleword element 0 of VSR[XT] as a half-precision value.

The contents of the leftmost 3 halfwords of doubleword element 0 of VSR[ XT] are set to 0 .
The contents of doubleword element 1 of VSR[ XT] are undefined.
3. \(F R\) is set to indicate if the rounded result was incremented. \(F \mid\) is set to 1 . FPRF is set to indicate the class and sign of the result.

For VSX Vector Floating-Point Arithmetic \({ }^{[1]}\) instructions, VSX Vector Floating-Point Reciprocal Estimate \({ }^{[2]}\) instructions, VSX Vector round and Convert Double-Precision to Single-Precision format (xvcvdpsp), VSX Vector Double-Precision to Integer Format Conversion \({ }^{[3]}\) instructions, and VSX Vector Integer to Floating-Point Format Conversion \({ }^{[4]}\) instructions, do the following.
1. \(X X\) is set to 1 .
2. Update of \(\mathrm{VSR}[\mathrm{XT}]\) is suppressed for all vector elements.
3. FR, FI, and FPRF are not modified.

For VSX Vector Convert Single-Precision to Half-Precision with round (xvcvsphp), do the following.
1. \(X X\) is set to 1 .
2. VSR[ \(X T]\) is not modified.
3. \(F R, F I\), and \(F P R F\) are not modified.
1. VSX Vector Floating-Point Arithmetic instructions: xvadddp, xvdivdp, xvmuldp, xvsubdp, xsaddsp, xvdivsp, xvmulsp, xvsubsp, xvmaddadp, xvmaddmdp, xvmsubadp, xvmsubmdp, xvnmaddadp, xvnmaddmdp, xvnmsubadp, xvnmsubmdp, xvmaddasp, xvmaddmsp, xvmsubasp, xvmsubmsp, xvnmaddasp, xvnmaddmsp, xvnmsubasp, xvnmsubmsp
2. VSX Vector Floating-Point Reciprocal Estimate instructions: xvredp, xvresp
3. VSX Vector Double-Precision to Integer Format Conversion instructions: xvcvdpsxds, xvcvdpsxws, xvcvdpuxds, xvcvdpuxws
4. VSX Vector Integer to Floating-Point Format Conversion instructions: xvcvsxddp, xvcvuxddp, xvcvsxdsp, xvcvuxdsp, xvcvsxwsp, xvcvuxwsp

\section*{418}

\subsection*{7.4.5.3 Action for \(X E=0\)}

When Inexact exception is disabled ( \(\mathrm{XE}=0\) ) and an Inexact exception occurs, the following actions are taken:
For VSX Scalar round and Convert Double-Precision to Single-Precision format (xscvdpsp), do the following.
1. \(X X\) is set to 1 .
2. The result is placed into word element 0 of \(\operatorname{VSR}[X T]\) as a single-precision value. The contents of word elements 1-3 of VSR[XT] are undefined.
3. FPRF is set to indicate the class and sign of the result.

For VSX Scalar Double-Precision Arithmetic \({ }^{[1]}\) instructions, VSX Scalar Single-Precision Arithmetic \({ }^{[2]}\) instructions, VSX Scalar Round to Single-Precision (xsrsp), the VSX Scalar Round to Double-Precision Integer Exact using Current rounding mode (xsrdpic), and VSX Scalar Integer to Double-Precision Format Conversion \({ }^{[3]}\) instructions, do the following.
1. \(X X\) is set to 1 .
2. The result is placed into doubleword element 0 of \(\operatorname{VSR}[\mathrm{XT}]\) as a double-precision value. The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.
3. FPRF is set to indicate the class and sign of the result.

For VSX Scalar Convert Double-Precision To Integer Word format with Saturate instructions,
xscvdpsxws, xscvdpuxws
do the following.
1. \(X X\) is set to 1 .
2. The result is placed into word element 1 of \(\operatorname{VSR}[X T]\). The contents of word elements 0,2 , and 3 of \(\mathrm{VSR}[\mathrm{XT}]\) are undefined.
3. FPRF is set to indicate the class and sign of the result.

For VSX Scalar round \& Convert Quad-Precision to Double-Precision (xscvqpdp), do the following.
1. \(X X\) is set to 1 .
2. The result is placed into the rightmost halfword of doubleword element 0 of VSR[XT] as a half-precision value.

The contents of the leftmost 3 halfwords of doubleword element 0 of VSR[ XT] are set to 0 .
The contents of doubleword element 1 of VSR[ XT] are undefined.
3. \(F R\) is set to indicate if the rounded result was incremented. \(F I\) is set to 1. FPRF is set to indicate the class and sign of the result.

\footnotetext{
1. VSX Scalar Double-Precision Arithmetic instructions:
xsadddp, xssubdp, xsmuldp, xsdivdp, xssqrtdp, xsmaddadp, xsmaddmdp, xsmsubadp, xsmsubmdp, xsnmaddadp, xsnmaddmdp, xsnmsubadp, xsnmsubmdp
2. VSX Scalar Single-Precision Arithmetic instructions:
xsaddsp, xssubsp, xsmulsp, xsdivsp, xssqrtsp, xsmaddasp, xsmaddmsp, xsmsubasp, xsmsubmsp, xsnmaddasp, xsnmaddmsp, xsnmsubasp, xsnmsubmsp
3. VSX Scalar Integer to Double-Precision Format Conversion instructions: xscvsxddp, xscvuxddp
}

For VSX Vector Double-Precision Arithmetic instructions,
xvadddp, xvsubdp, xvmuldp, xvdivdp, xvsqrtdp, xvmaddadp, xvmaddmdp, xvmsubadp, xvmsubmdp, xvnmaddadp, xvnmaddmdp, xvnmsubadp, xvnmsubmdp
do the following.
1. \(X X\) is set to 1 .
2. For each vector element causing an Inexact exception, the result is placed into its respective doubleword element of VSR[ XT] in double-precision format.
3. \(F R, F I\), and \(F P R F\) are not modified.

For any of the following instructions,
VSX Scalar Quad-Precision Arithmetic instructions:
xsaddqp[o], xsdivqp[o], xsmulqp[o], xssqrtqp[o], xssubqp[o]
xsmaddqp[o], xsmsubqp[o], xsnmaddqp[o], xsnmsubqp[o]
VSX Scalar Round Quad-Precision to Double-Extended-Precision (xsrqpxp)
VSX Scalar Round to Quad-Precision Integer (xsrqpi)
do the following.
1. \(X X\) is set to 1 .
2. The result is placed into VSR[ VRT +32 ] in quad-precision format.
3. \(F R\) is set to indicate if the rounded result was incremented. \(F I\) is set to 1 . \(F P R F\) is set to indicate the class and sign of the result.

For VSX Scalar round \& Convert Quad-Precision to Double-Precision (xscvqpdp), do the following.
1. \(X X\) is set to 1 .
2. The result is placed into doubleword element 0 of VSR[ VRT +32\(]\) in double-precision format.
\(0 \times 00000^{0} 0000_{2} 0000_{-} 0000\) is placed into doubleword element 1 of VSR[VRT+32].
3. \(F R\) is set to indicate if the rounded result was incremented. \(F I\) is set to 1 . \(F P R F\) is set to indicate the class and sign of the result.

For any of the following instructions,
VSX Scalar truncate \& Convert Quad-Precision to Signed Doubleword (xscvqpsdz) VSX Scalar truncate \& Convert Quad-Precision to Signed Word (xscvqpswz)
do the following.
1. \(X X\) is set to 1 .
2. The result is placed into doubleword element 0 of VSR[ VRT +32 ] in signed integer format. \(0 \times 0000_{-} 00000_{-} 0000_{2} 0000\) is placed into doubleword element 1 of VSR[VRT+32].
3. \(F R\) is set to \(0 . F I\) is set to 1.FPRF is undefined.

For any of the following instructions,
VSX Scalar truncate \& Convert Quad-Precision to Unsigned Doubleword (xscvqpudz) VSX Scalar truncate \& Convert Quad-Precision to Unsigned Word (xscvqpuwz)
do the following.
1. \(X X\) is set to 1 .
2. The result is placed into doubleword element 0 of VSR[ VRT +32 ] in unsigned integer format. \(0 \times 0000 \_0000 \_0000 \_0000\) is placed into doubleword element 1 of VSR[ VRT +32 ].
3. \(F R\) is set to \(0 . F I\) is set to 1 . \(F P R F\) is undefined.

For VSX Vector Convert Single-Precision to Half-Precision with round (xvcvsphp), do the following.
1. \(X X\) is set to 1 .
2. For each vector element causing an Underflow exception, the result is placed into the rightmost halfword of its respective word element of VSR[ XT] in half-precision format.

The contents of the leftmost halfword of its respective word element of VSR[ XT] are set to 0 .
3. \(F R, F I\), and \(F P R F\) are not modified.

For VSX Vector Single-Precision Arithmetic \({ }^{[1]}\) instructions, do the following.
1. \(X X\) is set to 1 .
2. For each vector element causing an Inexact exception, the result is placed into its respective word element of VSR[ XT] in single-precision format.
3. \(F R, F I\), and FPRF are not modified.

\footnotetext{
1. VSX Vector Single-Precision Arithmetic instructions:
xvaddsp, xvsubsp, xvmulsp, xvdivsp, xvsqrtsp, xvmaddasp, xvmaddmsp, xvmsubasp, xvmsubmsp, xvnmaddasp, xvnmaddmsp, xvnmsubasp, xvnmsubmsp
}

\subsection*{7.5 VSX Storage Access Operations}

The VSX Storage Access instructions compute the effective address (EA) of the storage to be accessed as described in Power ISA Book I.

\subsection*{7.5.1 Accessing Aligned Storage Operands}

The following quadword-aligned array, AH , consists of 8 halfwords.
```

short AW[4] = { 0x0001_0203,
0x0405_0607,
0x0809_0A0B,
OxOCOD_OEOF };

```

Figure 123 illustrates the Big-Endian storage image of array AW.


Figure 123.Big-Endian storage image of array AW
Figure 124 illustrates the Little-Endian storage image of array AW.


Figure 124.Little-Endian storage image of array AW
Figure 125 shows the result of loading that quadword into a VSR or, equivalently, shows the contents that must be in a VSR if storing that VSR is to produce the storage contents shown in Figure 123 for Big-Endian. Note that Figure shows the effect of loading the quadword from both Big-Endian storage and Little-Endian storage.


Figure 125.Vector-Scalar Register contents for aligned quadword Load or Store VSX Vector

\subsection*{7.5.2 Accessing Unaligned Storage Operands}

The following array, B, consists of 5 word elements.
```

int B[5]
B[0] = 0x01234567;
B[1] = 0x00112233
B[2] = 0x44556677;
B[3] = 0x8899AABB
B[4] = OxCCDDEEFF

```

Figure 126 illustrates both Big-Endian and Little-Endian storage images of array B.


Figure 126.Storage images of array B
Though this example shows the array starting at a quadword-aligned address, if the subject data of interest are elements 1 through 4, accessing elements 1 through 4 of array B involves an unaligned quadword storage access that spans two aligned quadwords.

\section*{Loading an Unaligned Quadword from Big-Endian Storage}

Loading elements from elements 1 through 4 of \(B\) (see Figure 126) into VR[VT] involves an unaligned quadword storage access.

VSX supports word-aligned vector and scalar storage accesses using Big-Endian byte ordering.

Big-Endian storage image of array B

\# Assumptions
GPR[Ra] = address of \(B\)
GPR[Rb] = 4 (index to \(\mathrm{B}[1]\) )
lxvw4x Xt,Ra,Rb

Figure 127.Process to load unaligned quadword from Big-Endian storage using Load VSX Vector Word*4 Indexed

\section*{Loading an Unaligned Quadword from Little-Endian Storage}

Loading elements from elements 1 through 4 of \(B\) (see Figure 126) into VR[VT] involves an unaligned quadword storage access.

VSX supports word-aligned vector and scalar storage accesses using Little-Endian byte ordering.

\# Assumptions
GPR[A] = address of \(B\)
\(\operatorname{GPR}[B]=4 \quad\) (index to \(\mathrm{B}[1]\) )
lxvw4x Xt,Ra,Rb
Xt:


Figure 128.Process to load unaligned quadword from Little-Endian storage Load VSX Vector Word*4 Indexed

\section*{Storing an Unaligned Quadword to Big-Endian Storage}

Storing a VSR to elements 1 through 4 of \(B\) (see Figure 126) into \(\mathrm{VR}[\mathrm{VT}]\) involves an unaligned quadword storage access.

VSX supports word-aligned vector and scalar storage accesses using Big-Endian byte ordering.

Big-Endian storage image of array \(B\)


\# Assumptions
GPR[Ra] = address of B
GPR[Rb] \(=4 \quad(\) index to \(B[1])\)
stxum4x Xs, Ra, Rb
0×0000:
0×0010:


Figure 129.Process to store unaligned quadword to Big-Endian storage using Store VSX Vector Word*4 Indexed

\section*{Storing an Unaligned Quadword to Little-Endian Storage}

Storing a VSR to elements 1 through 4 of \(B\) (see Figure 126) into VR[VT] involves an unaligned quadword storage access.

VSX supports word-aligned vector and scalar storage accesses using Little-Endian byte ordering.

Little-Endian storage image of array B


\# Assumptions
GPR[A] = address of B
GPR[B] \(=4 \quad(\) index to \(B[1])\)
stxvw4x Xs,Ra,Rb

Figure 130.Process to store unaligned quadword to Little-Endian storage Store VSX Vector Word*4 Indexed

\subsection*{7.5.3 Storage Access Exceptions}

Storage accesses cause the system data storage error handler to be invoked if the program is not allowed to modify the target storage (Store only), or if the program attempts to access storage that is unavailable.

\subsection*{7.6 VSX Instruction Set}

\subsection*{7.6.1 VSX Instruction Set Summary}

\subsection*{7.6.1.1 VSX Storage Access Instructions}

There are two basic forms of scalar load and scalar store instructions, word and doubleword. VSX Scalar Load instructions place a copy of the contents of the addressed word or doubleword in storage into the left-most word or doubleword element of the target VSR. The contents of the right-most element(s) of the target VSR are undefined. VSX Scalar Store instructions place a copy of the contents of the left-most word or doubleword element in the source VSR into the addressed word or doubleword in storage.

There are two basic forms of vector load and vector store instructions, a vector of 4 word elements and a vector of two doublewords. Both forms access a quadword in storage.

There is one basic form of vector load and splat instruction, doubleword. VSX Vector Load and Splat instruction places a copy of the contents of the addressed doubleword in storage into both doubleword elements of the target VSR.

\subsection*{7.6.1.1.1 VSX Scalar Storage Access Instructions}
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline Ixsdx & Load VSX Scalar Doubleword Indexed & 481 \\
\hline Ixsspx & Load VSX Scalar Single-Precision Indexed & 486 \\
\hline Ixsiwax & Load VSX Scalar as Integer Word Algebraic Indexed & 484 \\
\hline Ixsiwzx & Load VSX Scalar as Integer Word and Zero Indexed & 485 \\
\hline
\end{tabular}

Table 8. VSX Scalar Load Instructions
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline stxsdx & Store VSX Scalar Doubleword Indexed & 499 \\
\hline stxsspx & Store VSX Scalar Single-Precision Indexed & 503 \\
\hline stxsiwx & Store VSX Scalar as Integer Word Indexed & 501 \\
\hline
\end{tabular}

Table 9. VSX Scalar Store Instructions

\subsection*{7.6.1.1.2 VSX Vector Storage Access Instructions}
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline Ixvd2x & Load VSX Vector Doubleword*2 Indexed & 489 \\
\hline Ixvw4x & Load VSX Vector Word*4 Indexed & 497 \\
\hline
\end{tabular}

Table 10.VSX Vector Load Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline Ixvdsx & Load VSX Vector Doubleword and Splat Indexed & 495 \\
\hline
\end{tabular}

Table 11.VSX Vector Load and Splat Instruction
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline stxvd2x & Store VSX Vector Doubleword*2 Indexed & 505 \\
\hline stxvw4x & Store VSX Vector Word*4 Indexed & 507
\end{tabular}

Table 12.VSX Vector Store Instructions

\subsection*{7.6.1.2 VSX Move Instructions}

\subsection*{7.6.1.2.1 VSX Scalar Move Instructions}
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xsabsdp & VSX Scalar Absolute Value Double-Precision & 513 \\
\hline xscpsgndp & VSX Scalar Copy Sign Double-Precision & 535 \\
\hline xsnabsdp & VSX Scalar Negative Absolute Value Double-Precision & 608 \\
\hline xsnegdp & VSX Scalar Negate Double-Precision & 609 \\
\hline
\end{tabular}

Table 13.VSX Scalar Double-Precision Move Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xsabsqp & VSX Scalar Absolute Quad-Precision & 513 \\
\hline xsnegqp & VSX Scalar Negate Quad-Precision & 609 \\
\hline xsnabsqp & VSX Scalar Negative Absolute Quad-Precision & 608 \\
\hline
\end{tabular}

Table 14. VSX Scalar Quad-Precision Move Instructions

\subsection*{7.6.1.2.2 VSX Vector Move Instructions}
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvabsdp & VSX Vector Absolute Value Double-Precision & 660 \\
\hline xvcpsgndp & VSX Vector Copy Sign Double-Precision & 675 \\
\hline xvnabsdp & VSX Vector Negative Absolute Value Double-Precision & 729 \\
\hline xvnegdp & VSX Vector Negate Double-Precision & 730 \\
\hline
\end{tabular}

Table 15.VSX Vector Double-Precision Move Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvabssp & VSX Vector Absolute Value Single-Precision & 660 \\
\hline xvcpsgnsp & VSX Vector Copy Sign Single-Precision & 675 \\
\hline xvnabssp & VSX Vector Negative Absolute Value Single-Precision & 729 \\
\hline xvnegsp & VSX Vector Negate Single-Precision & 730 \\
\hline
\end{tabular}

Table 16.VSX Vector Single-Precision Move Instructions

\subsection*{7.6.1.3 VSX Floating-Point Arithmetic Instructions}

\subsection*{7.6.1.3.1 VSX Scalar Floating-Point Arithmetic Instructions}
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline\(x\) xadddp & VSX Scalar Add Double-Precision & 514 \\
\hline xsdivdp & VSX Scalar Divide Double-Precision & 564 \\
\hline\(x s m u l d p\) & VSX Scalar Multiply Double-Precision & 602 \\
\hline xsredp & VSX Scalar Reciprocal Estimate Double-Precision & 634 \\
\hline xsrsqrtedp & VSX Scalar Reciprocal Square Root Estimate Double-Precision & 641 \\
\hline xssqrtdp & VSX Scalar Square Root Double-Precision & 643 \\
\hline xssubdp & VSX Scalar Subtract Double-Precision & 647 \\
\hline xstdivdp & VSX Scalar Test for software Divide Double-Precision & 653 \\
\hline xstsqrtdp & VSX Scalar Test for software Square Root Double-Precision & 654 \\
\hline
\end{tabular}

Table 17.VSX Scalar Double-Precision Elementary Arithmetic Instructions
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xsaddsp & VSX Scalar Add Single-Precision & 519 \\
\hline xsdivsp & VSX Scalar Divide Single-Precision & 568 \\
\hline xsmulsp & VSX Scalar Multiply Single-Precision & 606 \\
\hline xsresp & VSX Scalar Reciprocal Estimate Single-Precision & 635 \\
\hline xsrsqrtesp & VSX Scalar Reciprocal Square Root Estimate Single-Precision & 642 \\
\hline xssqrtsp & VSX Scalar Square Root Single-Precision & 646 \\
\hline xssubsp & VSX Scalar Subtract Single-Precision & 651 \\
\hline
\end{tabular}

Table 18.VSX Scalar Single-Precision Elementary Arithmetic Instructions
\begin{tabular}{llc} 
Mnemonic & Instruction Name & Page \\
\hline xsaddqp[o] & VSX Scalar Add Quad-Precision [using round to Odd] & 521 \\
\hline xsdivqp[o] & VSX Scalar Divide Quad-Precision [using round to Odd] & 566 \\
\hline xsmulqp[o] & VSX Scalar Multiply Quad-Precision [using round to Odd] & 604 \\
\hline xsrqpxp & VSX Scalar Round Quad-Precision to Double-Extended-Precision & 638 \\
\hline xssqrtqp[o] & VSX Scalar Square Root Quad-Precision [using round to Odd] & 644 \\
\hline xssubqp[o] & VSX Scalar Subtract Quad-Precision [using round to Odd] & 649 \\
\hline
\end{tabular}
| Table 19. VSX Scalar Quad-Precision Elementary Arithmetic Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xsmaddadp & VSX Scalar Multiply-Add Type-A Double-Precision & 572 \\
\hline xsmaddmdp & VSX Scalar Multiply-Add Type-M Double-Precision & 572 \\
\hline xsmsubadp & VSX Scalar Multiply-Subtract Type-A Double-Precision & 593 \\
\hline xsmsubmdp & VSX Scalar Multiply-Subtract Type-M Double-Precision & 593 \\
\hline xsnmaddadp & VSX Scalar Negative Multiply-Add Type-A Double-Precision & 610 \\
\hline xsnmaddmdp & VSX Scalar Negative Multiply-Add Type-M Double-Precision & 610 \\
\hline xsnmsubadp & VSX Scalar Negative Multiply-Subtract Type-A Double-Precision & 621 \\
\hline xsnmsubmdp & VSX Scalar Negative Multiply-Subtract Type-M Double-Precision & 621 \\
\hline
\end{tabular}

Table 20.VSX Scalar Double-Precision Multiply-Add Arithmetic Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xsmaddasp & VSX Scalar Multiply-Add Type-A Single-Precision & 575 \\
\hline xsmaddmsp & VSX Scalar Multiply-Add Type-M Single-Precision & 575 \\
\hline xsmsubasp & VSX Scalar Multiply-Subtract Type-A Single-Precision & 596 \\
\hline xsmsubmsp & VSX Scalar Multiply-Subtract Type-M Single-Precision & 596 \\
\hline xsnmaddasp & VSX Scalar Negative Multiply-Add Type-A Single-Precision & 615 \\
\hline xsnmaddmsp & VSX Scalar Negative Multiply-Add Type-M Single-Precision & 615 \\
\hline xsnmsubasp & VSX Scalar Negative Multiply-Subtract Type-A Single-Precision & 624 \\
\hline xsnmsubmsp & VSX Scalar Negative Multiply-Subtract Type-M Single-Precision & 624 \\
\hline
\end{tabular}

Table 21.VSX Scalar Single-Precision Multiply-Add Arithmetic Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xsmaddqp[o] & VSX Scalar Multiply-Add Quad-Precision [using round to Odd] & 578 \\
\hline xsmsubqp[o] & VSX Scalar Multiply-Subtract Quad-Precision [using round to Odd] & 599 \\
\hline xsnmaddqp[o] & VSX Scalar Negative Multiply-Add Quad-Precision [using round to Odd] & 618 \\
\hline xsnmsubqp[o] & VSX Scalar Negative Multiply-Subtract Quad-Precision [using round to Odd] & 627 \\
\hline
\end{tabular}
\| Table 22. VSX Scalar Quad-Precision Multiply-Add Arithmetic Instructions

\subsection*{7.6.1.3.2 VSX Vector Floating-Point Arithmetic Instructions}
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvadddp & VSX Vector Add Double-Precision & 661 \\
\hline xvdivdp & VSX Vector Divide Double-Precision & 700 \\
\hline xvmuldp & VSX Vector Multiply Double-Precision & 725 \\
\hline xvredp & VSX Vector Reciprocal Estimate Double-Precision & 748 \\
\hline xvrsqrtedp & VSX Vector Reciprocal Square Root Estimate Double-Precision & 752 \\
\hline xvsqrtdp & VSX Vector Square Root Double-Precision & 755 \\
\hline xvsubdp & VSX Vector Subtract Double-Precision & 757 \\
\hline xvtdivdp & VSX Vector Test for software Divide Double-Precision & 761 \\
\hline xvtsqrtdp & VSX Vector Test for software Square Root Double-Precision & 763 \\
\hline
\end{tabular}

Table 23.VSX Vector Double-Precision Elementary Arithmetic Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvaddsp & VSX Vector Add Single-Precision & 665 \\
\hline xvdivsp & VSX Vector Divide Single-Precision & 702 \\
\hline xvmulsp & VSX Vector Multiply Single-Precision & 727 \\
\hline xvresp & VSX Vector Reciprocal Estimate Single-Precision & 749 \\
\hline xvrsqrtesp & VSX Vector Reciprocal Square Root Estimate Single-Precision & 754 \\
\hline xvsqrtsp & VSX Vector Square Root Single-Precision & 756 \\
\hline xvsubsp & VSX Vector Subtract Single-Precision & 759 \\
\hline xvtdivsp & VSX Vector Test for software Divide Single-Precision & 762 \\
\hline xvtsqrtsp & VSX Vector Test for software Square Root Single-Precision & 763 \\
\hline
\end{tabular}

Table 24.VSX Vector Single-Precision Elementary Arithmetic Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvmaddadp & VSX Vector Multiply-Add Type-A Double-Precision & 705 \\
\hline xvmaddmdp & VSX Vector Multiply-Add Type-M Double-Precision & 705 \\
\hline xvmsubadp & VSX Vector Multiply-Subtract Type-A Double-Precision & 719 \\
\hline xvmsubmdp & VSX Vector Multiply-Subtract Type-M Double-Precision & 719 \\
\hline xvnmaddadp & VSX Vector Negative Multiply-Add Type-A Double-Precision & 731 \\
\hline xvnmaddmdp & VSX Vector Negative Multiply-Add Type-M Double-Precision & 731 \\
\hline xvnmsubadp & VSX Vector Negative Multiply-Subtract Type-A Double-Precision & 739 \\
\hline xvnmsubmdp & VSX Vector Negative Multiply-Subtract Type-M Double-Precision & 739 \\
\hline
\end{tabular}

Table 25.VSX Vector Double-Precision Multiply-Add Arithmetic Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvmaddasp & VSX Vector Multiply-Add Type-A Single-Precision & 708 \\
\hline xvmaddmsp & VSX Vector Multiply-Add Type-M Single-Precision & 708 \\
\hline xvmsubasp & VSX Vector Multiply-Subtract Type-A Single-Precision & 722 \\
\hline xvmsubmsp & VSX Vector Multiply-Subtract Type-M Single-Precision & 722 \\
\hline xvnmaddasp & VSX Vector Negative Multiply-Add Type-A Single-Precision & 736 \\
\hline xvnmaddmsp & VSX Vector Negative Multiply-Add Type-M Single-Precision & 736 \\
\hline xvnmsubasp & VSX Vector Negative Multiply-Subtract Type-A Single-Precision & 742 \\
\hline xvnmsubmsp & VSX Vector Negative Multiply-Subtract Type-M Single-Precision & 742 \\
\hline
\end{tabular}

Table 26.VSX Vector Single-Precision Multiply-Add Arithmetic Instructions

\subsection*{7.6.1.4 VSX Floating-Point Compare Instructions}

\subsection*{7.6.1.4.1 VSX Scalar Floating-Point Compare Instructions}
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xscmpodp & VSX Scalar Compare Ordered Double-Precision & 529 \\
\hline xscmpudp & VSX Scalar Compare Unordered Double-Precision & 532 \\
\hline
\end{tabular}

Table 27.VSX Scalar Compare Double-Precision Instructions
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xsmaxdp & VSX Scalar Maximum Double-Precision & 581 \\
\hline xsmindp & VSX Scalar Minimum Double-Precision & 587 \\
\hline
\end{tabular}

Table 28.VSX Scalar Double-Precision Maximum/Minimum Instructions
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xscmpoqp & VSX Scalar Compare Ordered Quad-Precision & 531 \\
\hline xscmpuqp & VSX Scalar Compare Unordered Quad-Precision & 534 \\
\hline
\end{tabular}

I Table 29. VSX Scalar Quad-Precision Compare Instructions

\subsection*{7.6.1.4.2 VSX Vector Floating-Point Compare Instructions}
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvcmpeqdp[.] & VSX Vector Compare Equal To Double-Precision & 667 \\
\hline xvcmpgedp[.] & VSX Vector Compare Greater Than or Equal To Double-Precision & 669 \\
\hline xvcmpgtdp[.] & VSX Vector Compare Greater Than Double-Precision & 671 \\
\hline
\end{tabular}

Table 30.VSX Vector Compare Double-Precision Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvcmpeqsp[.] & VSX Vector Compare Equal To Single-Precision & 668 \\
\hline xvcmpgesp[.] & VSX Vector Compare Greater Than or Equal To Single-Precision & 670 \\
\hline xvcmpgtsp[.] & VSX Vector Compare Greater Than Single-Precision & 672 \\
\hline
\end{tabular}

Table 31.VSX Vector Compare Single-Precision Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvmaxdp & VSX Vector Maximum Double-Precision & 711 \\
\hline xvmindp & VSX Vector Minimum Double-Precision & 715 \\
\hline
\end{tabular}

Table 32.VSX Vector Double-Precision Maximum/Minimum Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvmaxsp & VSX Vector Maximum Single-Precision & 713 \\
\hline xvminsp & VSX Vector Minimum Single-Precision & 717 \\
\hline
\end{tabular}

Table 33.VSX Vector Single-Precision Maximum/Minimum Instructions

\subsection*{7.6.1.5 VSX FP-FP Conversion Instructions}
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xscvdpsp & VSX Scalar round and Convert Double-Precision to Single-Precision format & 538 \\
\hline xscvspdp & VSX Scalar Convert Single-Precision to Double-Precision format & 559 \\
\hline
\end{tabular}

Table 34.VSX Scalar DP-SP Conversion Instructions
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xscvqpdp[o] & VSX Scalar round \& Convert Quad-Precision to Double-Precision [using round to Odd] & 549 \\
\hline xscvdpqp & VSX Scalar Convert Double-Precision to Quad-Precision & \\
\hline
\end{tabular}

Table 35. VSX Scalar Quad-Precision Floating-Point Conversion Instructions
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xvcvdpsp & VSX Vector round and Convert Double-Precision to Single-Precision format & 676 \\
\hline xvcvspdp & VSX Vector Convert Single-Precision to Double-Precision format & 686 \\
\hline
\end{tabular}

Table 36.VSX Vector DP-SP Conversion Instructions

\subsection*{7.6.1.6 VSX FP-Integer Conversion Instructions}
7.6.1.6.1 VSX Scalar FP-Integer Conversion Instructions
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xscvdpsxds & \begin{tabular}{l} 
VSX Scalar truncate Double-Precision to integer and Convert to Signed Fixed-Point \\
Doubleword format with Saturate
\end{tabular} & 539 \\
\hline \multirow{2}{*}{ xscvdpsxws } & \begin{tabular}{l} 
VSX Scalar truncate Double-Precision to integer and Convert to Signed Fixed-Point Word \\
format with Saturate
\end{tabular} & 542 \\
\hline xscvdpuxds & \begin{tabular}{l} 
VSX Scalar truncate Double-Precision to integer and Convert to Unsigned Fixed-Point \\
Doubleword format with Saturate
\end{tabular} & 544 \\
\hline xscvdpuxws & \begin{tabular}{l} 
VSX Scalar truncate Double-Precision to integer and Convert to Unsigned Fixed-Point Word \\
format with Saturate
\end{tabular} & 546 \\
\hline
\end{tabular}

Table 37.VSX Scalar Double-Precision Convert to Integer Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xscvqpsdz & VSX Scalar truncate \& Convert Quad-Precision to Signed Dword & 550 \\
\hline xscvqpswz & VSX Scalar truncate \& Convert Quad-Precision to Signed Word & 552 \\
\hline xscvqpudz & VSX Scalar truncate \& Convert Quad-Precision to Unsigned Dword & 554 \\
\hline xscvqpuwz & VSX Scalar truncate \& Convert Quad-Precision to Unsigned Word & 556 \\
\hline
\end{tabular}

Table 38. VSX Scalar Quad-Precision Convert to Integer Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xscvsxddp & VSX Scalar Convert Signed Fixed-Point Doubleword to floating-point format and round to & 561 \\
\hline Double-Precision & & 563 \\
\hline
\end{tabular}

Table 39.VSX Scalar Double-Precision Convert from Integer Instructions
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xscvsdqp & VSX Scalar Convert Signed Dword to Quad-Precision & 558 \\
\hline xscvudqp & VSX Scalar Convert Unsigned Dword to Quad-Precision & 562 \\
\hline
\end{tabular}

Table 40. VSX Scalar Quad-Precision Convert from Integer Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xscvsxdsp & VSX Scalar Convert Signed Fixed-Point Doubleword to floating-point format and round to & 561 \\
\hline Single-Precision & \\
\hline xscvuxdsp & VSX Scalar Convert Unsigned Fixed-Point Doubleword to floating-point format and round to & 563 \\
\hline
\end{tabular}

Table 41.VSX Scalar Convert Integer to Single-Precision Instructions

\subsection*{7.6.1.6.2 VSX Vector FP-Integer Conversion Instructions}
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvcvdpsxds & \begin{tabular}{l} 
VSX Vector truncate Double-Precision to integer and Convert to Signed Fixed-Point \\
Doubleword format with Saturate
\end{tabular} & 677 \\
\hline xvcvdpsxws & \begin{tabular}{l} 
VSX Vector truncate Double-Precision to integer and Convert to Signed Fixed-Point Word \\
format with Saturate
\end{tabular} & 679 \\
\hline xvcvdpuxds & \begin{tabular}{l} 
VSX Vector truncate Double-Precision to integer and Convert to Unsigned Fixed-Point \\
Doubleword format with Saturate
\end{tabular} & 681 \\
\hline xvcvdpuxws & \begin{tabular}{l} 
VSX Vector truncate Double-Precision to integer and Convert to Unsigned Fixed-Point Word \\
format with Saturate
\end{tabular} & 683
\end{tabular}

Table 42.VSX Vector Convert Double-Precision to Integer Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvcvspsxds & \begin{tabular}{l} 
VSX Vector truncate Single-Precision to integer and Convert to Signed Fixed-Point \\
Doubleword format with Saturate
\end{tabular} & 688 \\
\hline xvcvspsxws & \begin{tabular}{l} 
VSX Vector truncate Single-Precision to integer and Convert to Signed Fixed-Point Word \\
format with Saturate
\end{tabular} & 690 \\
\hline xvcvspuxds & \begin{tabular}{l} 
VSX Vector truncate Single-Precision to integer and Convert to Unsigned Fixed-Point \\
Doubleword format with Saturate
\end{tabular} & 692 \\
\hline xvcvspuxws & \begin{tabular}{l} 
VSX Vector truncate Single-Precision to integer and Convert to Unsigned Fixed-Point Word \\
format with Saturate
\end{tabular} & 694
\end{tabular}

Table 43.VSX Vector Convert Single-Precision to Integer Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvcvsxddp & VSX Vector Convert and round Signed Fixed-Point Doubleword to Double-Precision format & 696 \\
\hline xvcvsxwdp & VSX Vector Convert Signed Fixed-Point Word to Double-Precision format & 697 \\
\hline xvcvuxddp & VSX Vector Convert and round Unsigned Fixed-Point Doubleword to Double-Precision format & 698 \\
\hline xvcvuxwdp & VSX Vector Convert Unsigned Fixed-Point Word to Double-Precision format & 699 \\
\hline
\end{tabular}

Table 44.VSX Vector Convert Integer to Double-Precision Instructions
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvcvsxdsp & VSX Vector Convert and round Signed Fixed-Point Doubleword to Single-Precision format & 696 \\
\hline xvcvsxwsp & VSX Vector Convert and round Signed Fixed-Point Word to Single-Precision format & 697 \\
\hline xvcvuxdsp & VSX Vector Convert and round Unsigned Fixed-Point Doubleword to Single-Precision format & 698 \\
\hline xvcvuxwsp & VSX Vector Convert and round Unsigned Fixed-Point Word to Single-Precision format & 699 \\
\hline
\end{tabular}

Table 45.VSX Vector Convert Integer to Single-Precision Instructions

\subsection*{7.6.1.7 VSX Round to Floating-Point Integer Instructions}

\subsection*{7.6.1.7.1 VSX Scalar Round to Floating-Point Integer Instructions}
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xsrdpi & VSX Scalar Round to Double-Precision Integer using round to Nearest Away & 630 \\
\hline xsrdpic & VSX Scalar Round to Double-Precision Integer Exact using Current rounding mode & 631 \\
\hline xsrdpim & VSX Scalar Round to Double-Precision Integer using round towards -Infinity rounding mode & 632 \\
\hline xsrdpip & VSX Scalar Round to Double-Precision Integer using round towards +Infinity rounding mode & 632 \\
\hline xsrdpiz & VSX Scalar Round to Double-Precision Integer using round towards Zero rounding mode & 633 \\
\hline
\end{tabular}

Table 46.VSX Scalar Round to Double-Precision Integer Instructions

\subsection*{7.6.1.7.2 VSX Vector Round to Floating-Point Integer Instructions}
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvrdpi & VSX Vector Round to Double-Precision Integer using round to Nearest Away & 745 \\
\hline xvrdpic & VSX Vector Round to Double-Precision Integer Exact using Current rounding mode & 745 \\
\hline xvrdpim & VSX Vector Round to Double-Precision Integer using round towards -Infinity rounding mode & 746 \\
\hline xvrdpip & VSX Vector Round to Double-Precision Integer using round towards +Infinity rounding mode & 746 \\
\hline xvrdpiz & VSX Vector Round to Double-Precision Integer using round towards Zero rounding mode & 747 \\
\hline
\end{tabular}

\section*{Table 47.VSX Vector Round to Double-Precision Integer Instructions}
\begin{tabular}{lll}
\hline Mnemonic & Instruction Name & Page \\
\hline\(x\) Pai
\end{tabular}
xsrqpi VSX Scalar Round to Quad-Precision Integer 636

Table 48. VSX Scalar Round to Quad-Precision Integer Instruction
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xvrspi & VSX Vector Round to Single-Precision Integer using round to Nearest Away & 750 \\
\hline xvrspic & VSX Vector Round to Single-Precision Integer Exact using Current rounding mode & 750 \\
\hline xvrspim & VSX Vector Round to Single-Precision Integer using round towards -Infinity rounding mode & 751 \\
\hline xvrspip & VSX Vector Round to Single-Precision Integer using round towards +Infinity rounding mode & 751 \\
\hline xvrspiz & VSX Vector Round to Single-Precision Integer using round towards Zero rounding mode & 752 \\
\hline
\end{tabular}

Table 49.VSX Vector Round to Single-Precision Integer Instructions

\subsection*{7.6.1.8 VSX Scalar Floating-Point Support Instructions}
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xscmpexpdp & VSX Scalar Compare Exponents Double-Precision & 523 \\
\hline xscmpexpqp & VSX Scalar Compare Exponents Quad-Precision & 524 \\
\hline xscpsgndp & VSX Scalar CopySign Double-Precision & 535 \\
\hline xscpsgnqp & VSX Scalar CopySign Quad-Precision & 535 \\
\hline xsiexpdp & VSX Scalar Insert Exponent Double-Precision & 570 \\
\hline xsiexpqp & VSX Scalar Insert Exponent Quad-Precision & 571 \\
\hline xststdcdp & VSX Scalar Test Data Class Double-Precision & 655 \\
\hline xststdcqp & VSX Scalar Test Data Class Quad-Precision & 656 \\
\hline xststdcsp & VSX Scalar Test Data Class Single-Precision & 657 \\
\hline xsxexpdp & VSX Scalar Extract Exponent Double-Precision & 658 \\
\hline xsxexpqp & VSX Scalar Extract Exponent Quad-Precision & 658 \\
\hline xsxigdp & VSX Scalar Extract Significand Double-Precision & 659 \\
\hline xsxsigqp & VSX Scalar Extract Significand Quad-Precision & 659 \\
\hline
\end{tabular}

Table 50. VSX Scalar Floating-Point Support Instructions

\subsection*{7.6.1.9 VSX Logical Instructions}
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xxland & VSX Logical AND & 771 \\
\hline xxlandc & VSX Logical AND with Complement & 771 \\
\hline xxlnor & VSX Logical NOR & 773 \\
\hline xxlor & VSX Logical OR & 774 \\
\hline xxlxor & VSX Logical XOR & 774 \\
\hline
\end{tabular}

\section*{Table 51.VSX Logical Instructions}
\begin{tabular}{llc} 
Mnemonic & Instruction Name & Page \\
\hline xxsel & VSX Select & 777 \\
\hline
\end{tabular}

Table 52.VSX Vector Select Instruction

\section*{Version 3.0}

\subsection*{7.6.1.10 VSX Permute Instructions}
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xxmrghw & VSX Merge High Word & 775 \\
\hline xxmrglw & VSX Merge Low Word & 775 \\
\hline
\end{tabular}

Table 53.VSX Merge Instructions
\begin{tabular}{llr}
\hline Mnemonic & Instruction Name & Page \\
\hline xxspltw & VSX Splat Word & 778 \\
\hline
\end{tabular}

Table 54.VSX Splat Instruction
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xxpermdi & VSX Permute Doubleword Immediate & 777 \\
\hline
\end{tabular}

Table 55.VSX Permute Instruction
\begin{tabular}{llc}
\hline Mnemonic & Instruction Name & Page \\
\hline xxsldwi & VSX Shift Left Double by Word Immediate & 778 \\
\hline
\end{tabular}

\section*{Table 56.VSX Shift Instruction}

\subsection*{7.6.2 VSX Instruction Description Conventions}

\subsection*{7.6.2.1 VSX Instruction RTL Operators}
x.bit[y]

Return the contents of bit \(y\) of \(x\).
x.bit[y:z]

Return the contents of bits \(y: z\) of \(x\).
x.word[y]

Return the contents of word element \(y\) of \(x\).

\section*{x.word[y:z]}

Return the contents of word elements \(y: z\) of \(x\).

\section*{x.dword[y]}

Return the contents of doubleword element \(y\) of \(x\).
x.dword[y:z]

Return the contents of doubleword elements \(y: z\) of \(x\).
\(x=y\)
The value of y is placed into x .
\(x \mathrm{l}=\mathrm{y}\)
The value of y is ORed with the value x and placed into \(x\).
~X
Return the one's complement of \(x\).
! \(\times\)
Return 1 if the contents of \(x\) are equal to 0 , otherwise return 0 .
\(\mathbf{x}|\mid \mathbf{y}\)
Return the value of \(x\) concatenated with the value of \(y\). For example, 0 b010 || Obll1 is the same as 0b010111.
\(x^{\wedge} y\)
Return the value of \(x\) exclusive ORed with the value of \(y\).
\(x\) ? y: z
If the value of \(x\) is true, return the value of \(y\), otherwise return the value \(z\).
\(\mathbf{x + y}\)
\(x\) and \(y\) are integer values.
Return the sum of \(x\) and \(y\).
\(\mathbf{x - y}\)
\(x\) and \(y\) are integer values.
Return the difference of \(x\) and \(y\).
\(\mathbf{x !}=\mathbf{y}\)
\(x\) and \(y\) are integer values.
Return 1 if x is not equal to y , otherwise return 0 .
\(\mathbf{x}<=\mathbf{y}\)
\(x\) and \(y\) are integer values.
Return 1 if \(x\) is less than or equal to \(y\), otherwise return 0 .
\(x>=y\)
\(x\) and \(y\) are integer values.
Return 1 if \(x\) is greater than or equal to \(y\), otherwise return 0 .

\subsection*{7.6.2.2 VSX Instruction RTL Function Calls}

\section*{AddDP( \(\mathbf{x}, \mathbf{y}\) )}
\(x\) and \(y\) are double-precision floating-point values.
If \(x\) or \(y\) is an \(\mathrm{SNaN}_{\text {, vxsnan_f }} \mathrm{f}\) ag is set to 1 .
If \(x\) is an Infinity and \(y\) is an Infinity of the opposite sign, vxisi_flag is set to 1 .
If x is a QNaN , return x .
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if \(y\) is a QNaN , return \(y\).
Otherwise, if y is an SNaN , return y represented as a QNaN.
Otherwise, if \(x\) and \(y\) are infinities of opposite sign, return the standard QNaN.
Otherwise, return the normalized sum of \(x\) and \(y\), having unbounded range and precision.

\section*{AddSP(x,y)}
\(x\) and \(y\) are single-precision floating-point values.
If \(x\) or \(y\) is an \(\mathrm{SNaN}_{\text {, vxsnan_f }} \mathrm{f}\) ag is set to 1 .
If \(x\) is an Infinity and \(y\) is an Infinity of the opposite sign, vxisi_flag is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if \(y\) is a QNaN , return \(y\).
Otherwise, if \(y\) is an SNaN , return y represented as a QNaN .
Otherwise, if \(x\) and \(y\) are infinities of opposite sign, return the standard QNaN.
Otherwise, return the normalized sum of \(x\) added to \(y\), having unbounded range and precision.
bf p_ABSOLUTE (x)
\(x\) is a binary floating-point value represented in the working floating-point format.
Return \(\times\) with sign set to 0 .
bf \(p_{-}\)ADD \((x, y)\)
\(x\) is a binary floating-point value represented in the working floating-point format.
\(y\) is a binary floating-point value represented in the working floating-point format.
If \(x\) or \(y\) is an \(S_{N a N, ~ v x s n a n \_f a g ~ i s ~ s e t ~ t o ~} 1\).
If \(x\) is an infinity and \(y\) is an infinity of the opposite sign, vxisi_flag is set to 1 .
If x is a QNaN , return x .
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN.
Otherwise, if \(y\) is a \(Q N a N\), return \(y\).
Otherwise, if \(y\) is an SNaN , return y represented as a QNaN.
Otherwise, if \(x\) and \(y\) are infinities of opposite sign, return the standard QNaN .
Otherwise, return the normalized sum of \(x\) and \(y\), having unbounded range and precision.
bf \(p_{-}\)COMPARE_EQ( \(\left.x, y\right)\)
\(x\) is a binary floating-point value represented in the working floating-point format.
\(y\) is a binary floating-point value represented in the working floating-point format.
Return \(0 b 0\) if \(x\) is NaN or y is a NaN .
Otherwise, return \(0 b 1\) if \(x\) is a Zero and \(y\) is a Zero.
Otherwise, return \(0 b 1\) if \(x\) is equal to \(y\).
Otherwise, return 0bo.
bf \(p_{-}\)COMPARE_GT(x, y)
\(x\) is a binary floating-point value represented in the working floating-point format.
\(y\) is a binary floating-point value represented in the working floating-point format.
Return \(0 b 0\) if \(x\) is NaN or y is a NaN .
Otherwise, return \(0 b 0\) if \(x\) is a Zero and \(y\) is a Zero.
Otherwise, return \(0 b 1\) if \(x\) is greater than \(y\).
Otherwise, return 0bo.
bfp_COMPARE_LT(x, y)
\(x\) is a binary floating-point value represented in the working floating-point format.
\(y\) is a binary floating-point value represented in the working floating-point format.
Return \(0 b 0\) if x is NaN or y is a NaN .
Otherwise, return \(0 b 0\) if \(x\) is a Zero and \(y\) is a Zero.
Otherwise, return \(0 b 1\) if \(x\) is less than \(y\).
Otherwise, return 0bo.
bf p_CONVERT_FROM_BFP16(x)
\(x\) is a floating-point value represented in half-precision format.
Let exponent be the contents of bits \(1: 5\) of \(x\).
Let \(f\) raction be the contents of bits 6:15 of \(x\).
Letresult.sign be set to 0 .
Let result. exponent be set to 0 .
Let result. significand be set to 0 .
Let result. class. 5 NaN be set to 0 .
Let result. class. QNaN be set to 0 .
Letresult.class.Infinity be set to 0 .
Let result. class. Zero be set to 0 .
Letresult.class. Denormal be set to 0 .
Letresult.class. Normal be set to 0 .
If \(x\) is a \(S N a N\), do the following.
result.class. SNaN is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
The contents of bit 0 of result. significand are set to 0 .
The contents of bits 1:10 of result. significand are set to the value of fraction.
Otherwise, if x is a QNaN, do the following.
result.class. QNaN is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
The contents of bit 0 of result. significand are set to 0 .
The contents of bits 1:10 of result. significand are set to the value of fraction.
Otherwise, if x is an Infinity value, do the following.
result.class.Infinity is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
Otherwise, if x is a Zero value, do the following.
result. class. Zero is set to 1.
result.sign is set to the contents of bit 0 of \(x\).

\section*{Version 3.0}

Otherwise, if x is a Denormal value, do the following.
result.class. Denormal is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
result.exp is set to the value - 14 .
The contents of bit 0 of result. significand are set to 0 .
The contents of bits 1:10 of result. significand are set to the value offraction.
result. significand is shifted left until the contents bit 0 of result. significand are equal to 1 .
result.exponent is decremented by the the number of bits result. significand was shifted.

\section*{Otherwise, do the following.}
result. class. Nor mal is set to 1.
result.sign is set to the contents of bit 0 of \(x\).
result. exp is set to the value of exponent subtracted by 15 .
The contents of bit 0 of result. significand are set to 1 .
The contents of bits 1:10 of result. significand are set to the value of fraction.
Return result.
bf P_CONVERT_FROM_BFP32(x)
\(x\) is a floating-point value represented in single-precision format.
Let exponent be the contents of bits 1:8 of \(x\).
Let \(f\) raction be the contents of bits 9:31 of \(x\).
Letresult. sign be initialized to 0 .
Let result. exponent be initialized to 0 .
Let result. significand be initialized to 0 .
Let result. class. SNaN be initialized to 0 .
Let result. class. QNaN be initialized to 0 .
Letresult. class. Infinity be initialized to 0 .
Let result. class. Zero be initialized to 0 .
Letresult. class. Denormal be initialized to 0 .
Letresult.class. Normal be initialized to 0 .
If x is a SNaN , do the following.
result.class. SNaN is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
The contents of bit 0 of result. significand are set to 0 .
The contents of bits 1:23 of result. significand are set to the value offraction.
Otherwise, if x is a QNaN , do the following.
result.class. QNaN is set to 1.
result.sign is set to the contents of bit 0 of \(x\).
The contents of bit 0 of result. significand are set to 0 .
The contents of bits 1:23 of result. significand are set to the value offraction.
Otherwise, if x is an Infinity value, do the following.
result.class.Infinity is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
Otherwise, if x is a Zero value, do the following.
result.class. Zero is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
Otherwise, if \(x\) is a Denormal value, do the following.
result.class. Denormal is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
result. exponent is set to the value - 126 .
The contents of bit 0 of result. significand are set to 0 .
The contents of bits \(1: 23\) of result. significand are set to the value of fraction.
result. significand is shifted left until the contents bit 0 of result. significand are equal to 1 .
result.exponent is decremented by the the number of bits result. significand was shifted.
Otherwise, do the following.
result.class. Normal is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
result. exponent is set to the value of exponent subtracted by 127.
The contents of bit 0 of result. significand are set to 1 .
The contents of bits \(1: 23\) of result. significand are set to the value of fraction.
Return result.
bf \(p_{-}\)CONVERT_FROM_BFP64(x)
\(x\) is a binary floating-point value represented in double-precision format.
Let exponent be the contents of bits \(1: 11\) of \(x\).
Let fracti on be the contents of bits 12:63 of \(x\).
result. sign is initialized to 0 .
result.exponent is initialized to 0 .
result. significand is initialized to 0 .
result. class. SNaN is initialized to 0 .
result.class. QNaN is initialized to 0 .
result.class. Infinity is initialized to 0 .
result. class. Zero is initialized to 0 .
result.class. Denormal is initialized to 0 .
result. class. Normal is initialized to 0 .
If x is a SNaN , do the following.
result. class. SNaN is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
The contents of bit 0 of result. significand are set to 0 .
The contents of bits 1:52 of result. significand are set to the value of fraction. The contents of the rest of result. significand are set to 0 .

Otherwise, if x is a QNaN , do the following.
result. class. QNaN is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
The contents of bit 0 of result. significand are set to 0 .
The contents of bits \(1: 52\) of result. significand are set to the value of fraction. The contents of the rest of result. significand are set to 0 .

Otherwise, if x is an Infinity, do the following.
result.class.Infinity is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
Otherwise, if x is a Zero, do the following.
result. class. Zero is set to 1.
result.sign is set to the contents of bit 0 of \(x\).
Otherwise, if x is a Denormal, do the following.
result.class. Denormal is set to 1 .
result. sign is set to the contents of bit 0 of \(x\).
result.exp is set to the value -1022 .
The contents of bit 0 of result. significand are set to 0 .
The contents of bits \(1: 52\) of result. significand are set to the value of fraction. The contents of the rest of result. significand are set to 0 .
result. significand is shifted left until the contents bit 0 of \(r\) esult. significand are equal to 1 .
result.exponent is decremented by the the number of bits result. significand was shifted.
Otherwise, do the following.
result.class. Normal is set to 1 .
result. sign is set to the contents of bit 0 of \(x\).
result.exp is set to the value of exponent subtracted by 1023 .
The contents of bit 0 of result. significand are set to 1 .
The contents of bits 1:52 of result. significand are set to the value of fraction.
The contents of the rest of result, significand are set to 0 .
Return result (i.e., the value \(x\) in the working floating-point format).

\section*{bf p_CONVERT_FROM_BFP128(x)}
x is a binary floating-point value represented in quad-precision format.
Let exponent be the contents of bits \(1: 15\) of \(x\).
Let \(f\) raction be the contents of bits 16:127 of \(x\).
result. sign is initialized to 0 .
result. exponent is initialized to 0 .
result. significand is initialized to 0 .
result. class. 5 NaN is initialized to 0 .
result. class. QNaN is initialized to 0 .
result. class. Infinity is initialized to 0 .
result. class. Zero is initialized to 0 .
result.class. Denormal is initialized to 0 .
result. class. Nor mal is initialized to 0 .
If \(x\) is a SNaN , do the following.
result. class. SNaN is set to 1 .
result. sign is set to the contents of bit 0 of \(x\).
The contents of bit 0 of result. significand are set to 0 .
The contents of bits 1:112 of result. significand are set to the value of fraction.
The contents of the rest of result. significand are set to 0 .
Otherwise, if x is a QNaN , do the following.
result. class. QNaN is set to 1.
result. sign is set to the contents of bit 0 of \(x\).
The contents of bit 0 of result. significand are set to 0 .
The contents of bits \(1: 112\) of result, significand are set to the value of fraction.
The contents of the rest of result. significand are set to 0 .
Otherwise, if x is an Infinity, do the following.
result.class. Infinity is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
Otherwise, if x is a Zero, do the following.
result. class. Zero is set to 1.
result.sign is set to the contents of bit 0 of \(x\).
Otherwise, if \(x\) is a Denormal, do the following.
result.class. Denormal is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
result.exp is set to the value -16382 .
The contents of bit 0 of result, significand are set to 0 .
The contents of bits 1:112 of result. significand are set to the value of fraction.
The contents of the rest of result. significand are set to 0 .
result.significand is shifted left until the contents bit 0 of result. significand are equal to 1 .
result. exponent is decremented by the the number of bits result. significand was shifted.
Otherwise, do the following.
result.class. Normal is set to 1 .
result.sign is set to the contents of bit 0 of \(x\).
result.exp is set to the value of exponent subtracted by 16383 .
The contents of bit 0 of result. significand are set to 1 .
The contents of bits 1:112 of result, significand are set to the value of fraction.
The contents of the rest of result. significand are set to 0 .
Return result (i.e., the value \(x\) in the working floating-point format).

\section*{Version 3.0}
bf \(p_{-}\)CONVERT_FROM_SI 64(x)
\(x\) is an integer value represented in signed doubleword integer format.
result.sign is initialized to 0 .
result.exponent is initialized to 0 .
result. significand is initialized to 0 .
result. class. 5 NaN is initialized to 0 .
result. class. QNaN is initialized to 0 .
result.class. Infinity is initialized to 0 .
result.class. Zero is initialized to 0 .
result.class. Denormal is initialized to 0 .
result.class. Normal is initialized to 0 .

If \(x\) is equal to \(0 \times 0000_{1} 0000 \_0000 \_0000\), result.class. Zero is set to 1.

Otherwise, do the following.
result.class. Normal is set to 1.
result.sign is set to the contents of bit 0 of \(x\).
result.exponent is set to the value 64.
Bits 0:64 of result. significand are set to the value of \(x\) sign-extended to 65 bits.
If bit 0 of result. significand is equal to 1 ,
result.sign is set to 1 , and
result. significand is set to the value of the two's complement of result. significand.
If bit 0 of result. significand is equal to 0 ,
result.significand is shifted left until bit 0 of result.significand is equal to 1 , and
result, exponent is decremented by the number of bitsresult, significand is shifted.
Return result (i.e., the value \(x\) in the working floating-point format).
bf p_CONVERT_FROM_UI 64(x)
\(x\) is an integer value represented in unsigned doubleword integer format.
Return x in the working floating-point format.
result. sign is initialized to 0 .
result. exponent is initialized to 0 .
result. significand is initialized to 0 .
result. class. 5 NaN is initialized to 0 .
result. class. QNaN is initialized to 0 .
result. class. Infinity is initialized to 0 .
result. class. Zero is initialized to 0 .
result. class. Denormal is initialized to 0 .
result. class. Nor mal is initialized to 0 .
If \(x\) is equal to \(0 \times 0000 \_0000 \_0000 \_0000\), do the following. result. class. Zero is set to \({ }^{-1}\).

Otherwise, do the following.
result.class. Normal is set to 1.
result. sign is set to 0 .
result. exponent is set to the value 64.
Bits 0:64 of result. significand is set to the value of \(x\) zero-extended to 65 bits.
If bit 0 of result.significand is equal to 0 , result.significand is shifted left until bit 0 of result.significand is equal to 1 and result.exponent is decremented by the number of bits result.significand is shifted.

Return result (i.e., the value x in the working floating-point format).
bf p_CONVERT_TO_BFP16(x)
\(x\) is a floating-point value represented in the working format.
If \(x\). cl ass. \(Q N a N=1\), do the following.
Bit 0 of result is set to the value of \(x\). sign.
Bits 1:5 of \(r\) esult are set to the value 0 b11111.
Bits 6:15 of result are set to the value of bits 1:10 of x. significand.
Otherwise, if \(x\). class. Infinity \(=1\), do the following.
Bit 0 of \(r\) esult is set to the value of \(x\). sign.
Bits 1:5 of \(r\) esult are set to the value 0 b11111.
Bits 6:15 of result are set to 0 .
Otherwise, if \(x\), class. Zer \(0=1\), do the following.
Bit 0 of result is set to the value of \(x\). sign.
Bits 1:15 of result are set to 0 .
Otherwise, if \(x\). exponent is less than -14 and \(U E=0\), do the following.
Bit 0 of \(r\) esult is set to the value of \(x\). sign.
sh_cnt is set to the difference, \(\cdot 14\) - x. exponent.
Bits 1:5 of result are set to \(0 b 00000\).
Bits \(6: 15\) of result are set to bits \(1: 10\) of \(x\). significand shifted right by sh_cnt bits.
Otherwise, if \(x\). exponent is less than -14 and \(U E=1\), result is undefined.
Otherwise, if \(x\). exponent is greater than 15 and \(0 E=1\), result is undefined.

\section*{Version 3.0}

Otherwise, do the following.
Bit 0 of result is set to the value of \(x\). sign.
Bits \(1: 5\) of result are set to the sum, \(x\). exponent +15 .
Bits 6:15 of result are set to bits 1:10 of x. significand.
Return result.
bfp_CONVERT_TO_BFP32(x)
\(x\) is a floating-point value represented in the working format.
If \(\mathrm{x}, \mathrm{cl}\) ass,\(~ Q N a N=1\), do the following.
Bit 0 of result is set to the value of \(x\). sign.
Bits \(1: 8\) of result are set to the value 0 b1111_1111.
Bits 9:31 of result are set to the value of bits 1:23 of x. si gnificand.
Otherwise, if \(x\). class. Infinity \(=1\), do the following.
Bit 0 of \(r\) esult is set to the value of \(x\). sign.
Bits 1:9 of result are set to the value ob1111_1111.
Bits \(9: 31\) of result are set to 0 .
Otherwise, if \(x, ~ c l a s s\). Zer \(0=1\), do the following.
Bit 0 of result is set to the value of \(x\). sign.
Bits \(1: 31\) of \(r\) esult are set to 0 .
Otherwise, if \(x\). exponent is less than 126 and \(U E=0\), do the following.
Bit 0 of \(r\) esult is set to the value of \(x\). sign.
sh_cnt is set to the difference, \(\cdot 126\) - x. exponent.
Bits 1:8 of result are set to \(0 b 0000 \_0000\).
Bits \(9: 31\) of result are set to bits \(1: 23\) of \(x\). si gnificand shifted right by sh_cnt bits.
Otherwise, if \(x\). exponent is less than -126 and \(U E=1\), result is undefined.
Otherwise, if \(x\). exponent is greater than 127 and \(O E=1\), result is undefined.
Otherwise, do the following.
Bit 0 of \(r\) esult is set to the value of \(x\). sign.
Bits 1:8 of result are set to the sum, \(x\). exponent +127 .
Bits 9:31 of \(r\) esult are set to bits 1:23 of \(x\). significand.
Returnresult.
bf p_CONVERT_TO_BFP64(x)
\(x\) is a floating-point value represented in the working format.
If \(x\), class. \(Q N a N=1\), do the following.
Bit 0 of result is set to the value of \(x\). sign.
Bits 1:11 of res ult are set to the value 0b111_1111_1111.
Bits 12:63 of result are set to the value of bits 1:52 of \(x\). significand.
Otherwise, if x . class. I nfinity \(=1\), do the following.
Bit 0 of result is set to the value of \(x\). sign.
Bits 1:11 of \(r\) esult are set to the value 0b111_1111_1111.
Bits 12:63 of result are set to 0 .
Otherwise, if \(x\). class. Zer \(0=1\), do the following.
Bit 0 of result is set to the value of \(x\). sign. Bits \(1: 63\) of result are set to 0 .

Otherwise, if \(x\). exponent is less than 1022 and \(U E=0\), do the following.
Bit 0 of \(r\) esult is set to the value of \(x\). sign.
sh_cnt is set to the difference, \(\cdot 1022\) - x. exponent.
Bits 1:11 of result are set to 0b000_0000_0000.
Bits 12:63 of result are set to bits \(\overline{1}: 52\) of x . significand shifted right by sh _ cnt bits.
Otherwise, if \(x\). exponent is less than -1022 and \(U E=1\), result is undefined.
Otherwise, if \(x\). exponent is greater than 1023 and \(0 E=1\), result is undefined.
Otherwise, do the following.
Bit 0 of result is set to the value of x . sign.
Bits \(1: 11\) of result are set to the sum, x. exponent +1023 .
Bits 12:63 of result are set to bits 1:52 of \(x\). significand.
Return result.
bf p_CONVERT_TO_BFP128(x)
\(x\) is a quad-precision floating-point value that is represented in the working floating-point format.
If \(x\) is a QNaN ,
the contents of bit 0 of \(r\) esult are set to the value of \(x\). sign,
the contents of bits 1:15 of result are set to the value 0b111_1111_1111_1111, and
the contents of bits \(16: 127\) of result are set to the value of bits \(1: 112\) of \(x\). significand.
Otherwise, if x is a Zero,
the contents of bit 0 of result are set to the value of \(x\). sign, and
the contents of bits \(1: 15\) of result are set to the value \(0 b 000 \_0000 \_0000 \_0000\), and
the contents of bits 16:127 of \(r\) esult are set to the value \(0 \times 0000 \_00000_{-}^{-0} 000 \_0000_{-} 0000 \_0000 \_0000\).
Otherwise, if \(x\) is an Infinity,
the contents of bit 0 of result are set to the value of \(x . \operatorname{sign}\),
the contents of bits \(1: 15\) of result are set to the value 0b111_1111_1111_1111, and
the contents of bits 16:127 of \(r\) esult are set to the value \(0 \times 0000 \_00000_{-}^{-0000} 0000 \_0000_{0} 0000_{-} 0000\).
Otherwise, do the following.
If the exponent of \(x\) is less than -16382 ,
the contents of bit 0 of result are set to the value of \(x\). sign,
the contents of bits \(1: 15\) of result are set to the value \(06000 \_0000 \_0000 \_0000\), and
the contents of bits 16:127 of result are set to the value of bits 1:112 of the significand of \(x\) shifted right by \(N\) bits, where \(N\) is the value 16382 subtracted by the value of the exponent of \(x\).
Otherwise,
the contents of bit 0 of result are set to the value of \(x\). sign,
the contents of bits 1:15 of result are set to the sum of the exponent of \(x\) and 16383, and the contents of bits \(16: 127\) of \(r\) esult are set to the value of bits \(1: 112\) of the significand of \(x\).

Return result (i.e., \(x\) in quad-precision format).
bf p_CONVERT_TO_SI 64(x)
\(x\) is an integer value represented in the working floating-point format.
Return the value \(x\) in signed doubleword integer format.
bf p_CONVERT_TO_U164(x)
\(x\) is an integer value represented in the working floating-point format.
Return the value x in 64-bit unsigned integer format.
bf \(p_{-}\)DENORM \((x, y)\)
\(x\) is an integer value specifying the target format's \(E\) min value.
\(y\) is a binary floating-point value that is represented in the working floating-point format.
If \(y\). exponent is less than Emin , letsh_cnt be the value Emin - y. exponent.
Otherwise, let \(s h_{\text {_ }}\) cnt be the value 0 .
y. significand, having unbounded precision, is shifted right by sh_cnt bits.
y. exponent is incremented bysh_cnt.

Return \(y\) in the working floating-point format.

\section*{bf p_DIVIDE( \(x, y)\)}
\(x\) is a binary floating-point value that is represented in the working floating-point format.
\(y\) is a binary floating-point value that is represented in the working floating-point format.
If \(x\) or \(y\) is an \(S N a N\), \(v x s n a n \_f l a g\) is set to 1 .
Otherwise, if \(x\) and \(y\) are infinities, vxidi \(f l a g\) is set to 1 .
Otherwise, if \(x\) and \(y\) are zeros, \(v x z d z_{-} f\) ag is set to 1 .
Otherwise, if \(x\) is a finite value and \(y\) is a zero, \(z x_{-} f \mid a g\) is set to 1 .
If x is a QNaN , return x .
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if \(y\) is a QNaN , return \(y\).
Otherwise, if y is an SNaN , return y represented as a QNaN .
Otherwise, if \(x\) and \(y\) are infinities, return the standard QNaN .
Otherwise, if \(x\) and \(y\) are zeros, return the standard QNaN.
Otherwise, if \(y\) is a zero, return infinity, having the sign of the exclusive-OR of the signs of \(x\) and \(y\).
Otherwise, return the normalized quotient of \(x \div y\), having unbounded range and precision.
bfp_INFINITY()
Return a positive floating-point infinity value, represented in the working format.
```

bfp_INITIALIZE(result)
result.class.Infinity }\leftarrow
return(result)

```
bf p_INITIALIZE(x)
Let x . sign be set to 0 .
Let \(x\). exponent be set to 0 .
Let x . significand be set to 0 .
Let x . class. SN N N be set to 0 .
Let x . class. QNaN be set to 0 .
Let x . class. Infinity be set to 0 .
Let x . class. Zero be set to 0 .
Let x . class. Denormal be set to 0 .
Let \(x\). class. Normal be set to 0 .
Return x .
bf P_MULTIPLY( \(x, y)\)
\(x\) is a binary floating-point value represented in the working floating-point format.
\(y\) is a binary floating-point value represented in the working floating-point format.
If \(x\) or \(y\) is an \(\mathrm{SNaN}_{\text {, vxsnan_ }} \mathrm{fl}\) ag is set to 1 .
Otherwise, if \(x\) is an infinity and \(y\) is a zero, \(v x i m z f l a g\) is set to 1 .
Otherwise, if \(x\) is a zero and \(y\) is an infinity, \(\mathrm{vximz}_{-}^{-} \mathrm{f} \mid \mathrm{ag}\) is set to 1 .
If x is a QNaN , return x .
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if \(y\) is a \(Q N a N\), return \(y\).
Otherwise, if \(y\) is an SNaN , return y represented as a QNaN.
Otherwise, if \(x\) is an infinity and \(y\) is a zero, return the standard QNaN .
Otherwise, if \(x\) is a zero and \(y\) is an infinity, return the standard \(Q N a N\).
Otherwise, return the normalized product of \(x \times y\), having unbounded range and precision.
bf p_MULTI PLY_ADD( \(x, y, z)\)
\(x\) is a binary floating-point value represented in the working floating-point format.
\(y\) is a binary floating-point value represented in the working floating-point format.
\(z\) is a binary floating-point value represented in the working floating-point format.
If \(x, y\), or \(z\) is an SNaN, vxsnan_flag is set to 1 .
Otherwise, if \(x\) is an infinity and \(y\) is a zero, \(v x i m z f l a g\) is set to 1 .
Otherwise, if \(x\) is a zero and \(y\) is an infinity, \(v x i \mathrm{mz}^{-} f l \mathrm{ag}\) is set to 1 .
Otherwise, if \(z\) and the product of \(x \times y\) are Infinity values having opposite signs, vxisi_flag is set to 1 .
If x is a QNaN , return x .
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if \(z\) is a QNaN , return \(z\).
Otherwise, if \(z\) is an SNaN , return \(z\) represented as a QNaN .
Otherwise, if \(y\) is a \(Q N a N\), return \(y\).
Otherwise, if \(y\) is an SNaN , return y represented as a QNaN.
Otherwise, if \(x\) is an infinity and \(y\) is a zero, return the standard \(Q N a N\).
Otherwise, if \(x\) is a zero and \(y\) is an infinity, return the standard QNaN.
Otherwise, if \(z\) and the product of \(x \times y\) are Infinity values having opposite signs, return the standard QNaN .
Otherwise, return the sum of \(z\) and the normalized product of \(x \times y\), having unbounded range and precision.
bf p_NEGATE ( \(x\) )
\(x\) is a binary floating-point value that is represented in the working floating-point format.
Return \(\times\) with its sign complemented.
bf p_NMAX_BFP16()
Return the largest, positive, normalized half-precision floating-point value, \(\left(2 \cdot 2^{\cdot 10}\right) \times 2^{+15}\), represented in the working format.
```

bfp_INITIALIZE(result)
result.exponent }\leftarrow+1
result.significand.bit[0:10] \leftarrowOb111_1111_1111
result.class.Normal }\leftarrow
return(result)

```
bf p_NMAX_BFP64
Return the largest finite double-precision value (i.e., \(2^{1024} .2^{1024-53}\) ) in the working floating-point format.
```

return(bfp_CONVERT_FROM_BFP64(Ox7FEF_FFFF_FFFF_FFFF))

```
bf p_NMAX_BFP80
Return the largest finite double-extended-precision value (i.e., \(2^{16384} \cdot 2^{16384 \cdot 65}\) ) in the working floating-point format.
```

return( bfp_CONVERT_FROM_BFP80(Ox7FFE_FFFF_FFFF_FFFF_FFFF))

```
bf P_NMAX_BFP128
Retürn the largest finite quad-precision value (i.e., \(2^{16384} .2^{16384-113}\) ) in the working floating-point format.
```

return( bf P_CONVERT_FROM_BFP128(Ox7FFE_FFFF_FFFF_FFFF_FFFF_FFFF_FFFF_FFFF) )

```
bf p_NMIN_BFP16()
Return the smallest, positive, normalized half-precision floating-point value, \(2^{\cdot 14}\), represented in the working format.
```

bfp_INITIALIZE(result)
result.exponent \leftarrow.14
result.significand.bit[0:10]}\leftarrow0b100_0000_0000
result.class.Normal }\leftarrow

```

\section*{return(result)}

\section*{bf p_NMIN_BFP64}

Retürn the smallest, positive, normalized double-precision value, \(2 \cdot 1022\), represented in the binary floating-point working format.
```

return( bfp_CONVERT_FROM_BFP64(0x0010_0000_0000_O000) )

```
bf p_NMIN_BFP80
Return the smallest, positive, normalized double-extended-precision value, \(2^{\cdot 16382}\), represented in the binary floating-point working format.
```

return( bfp_CONVERT_FROM_BFP80(0x0001_0000_0000_0000_0000) )

```
bf p_NMI N_BFP128
Return the smallest, positive, normalized quad-precision value, \(2^{-16382}\), represented in the binary floating-point working format.
```

return( bfp_CONVERT_FROM_BFP128(0x0001_0000_0000_0000_0000_0000_0000_0000) )

```
bf p_QUIET(x)
\(x\) is a Signalling NaN .
Return x converted to a Quiet NaN with \(\mathrm{x}, \mathrm{cl}\) ass. QNaN set to 1 and x . class. SN N N set to 0 .
bf p_ROUND_CEIL(p, x)
\(x\) is à binary floating-point value that is represented in the working floating-point format and has unbounded exponent range and significand precision. x must be rounded as presented, without prenormalization.
\(p\) is an integer value specifying the precision (i.e., number of bits) the significand is rounded to.
Return the smallest floating-point number having unbounded exponent range and a significand with a width of \(p\) bits that is greater or equal in value to \(x\).
inc_flag is set to 1 if the magnitude of the value returned is greater than \(x\).
\(x x_{-} \bar{f} \mathrm{lag}\) is set to 1 if the value returned is not equal to \(x\).
bf p_ROUND_FLOOR (p, x)
\(x\) is a binary floating-point value that is represented in the working floating-point format and has unbounded exponent range and significand precision. The value must be rounded as presented, without prenormalization.
\(p\) is an integer value specifying the precision (i.e., number of bits) the significand is rounded to.
Return the largest floating-point number having unbounded exponent range and a significand with a width of \(p\) bits that is lesser or equal in value to \(x\).
inc_flag is set to 1 if the magnitude of the value returned is greater than \(x\).
\(x x_{-} \bar{f} \mathrm{lag}\) is set to 1 if the value returned is not equal to \(x\).
bf P_ROUND_TO_BFP16(x,y)
\(y\) is a normalized floating-point value represented in the working format, having unbounded exponent range and significand precision.
\(x\) is a 2-bit integer value specifying one of four rounding modes.
\(0 b 00\) Round to Nearest Even
Ob01 Round towards Zero
Ob10 Round towards +Infinity
Ob11 Round towards - Infinity
If \(y\) is an QNaN, Infinity, or Zero, return \(y\). Otherwise, if \(y\) is an SNaN, set vxsnan_flag to 1 and return the corresponding QNaN representation of \(y\). Otherwise, return the value y rounded to half-precision format's exponent range and significand precision using the rounding mode specified by x .
```

if y.class.Zero | y.class.Infinity then return(y)
if y.class.QNaN | y.class.SNaN then do
result \leftarrowy
result.significand.bit[1] \leftarrow1
result.significand, bit[11:inf]}\leftarrow
result.class.SNaN \leftarrow0
result.class.QNaN \leftarrow1
vxsnan_flag \leftarrowy.class.SNaN
return(result)
end
if bfp_COMPARE_LT(y,bfp_NMIN_BFP16()) then do
if FPSCR.UE=0 then do
do while y.exponent <-14 || denormalize y
y.significand \leftarrowy.significand >> 1
y.exponent \leftarrowy.exponent + 1
end
if x=ObOO then result \& bf _ ROUND_TO_BFP16_NEAR_EVEN(y)
if x=0b01 then result \&bf __ROUND_TO_BFP16_TRUNC(y)
if x=Ob10 then result \& bf P_ROUND_TO_BFP16_CEIL(y)
if x=0b11 then result \leftarrowbf p_ROUND_TO_BFP16_FLOOR(y)
do while result.significand.bit[0] = 0 ll normalize result
result.significand < result.significand << 1
result.exponent \leftarrowresult.exponent - 1
end
ux_flag\& <x_flag
return(result)
end
else do
y.exponent \&y.exponent + 24
ux_flag \leftarrow1
end
end
if x=0bOO then result \leftarrow bfp_ROUND_TO_BFP16_NEAR_EVEN(y)
if x=0b01 then result \& bf __ROUND_TO_BFP16_TRUNC(y)
if x=0b10 then result \& bfp_ROUND_TO_BFP16_CEIL(y)
if x=0b11 then result \leftarrowbf __ROUND_TO_BFP16_FLOOR(y)

```
```

if bfp_COMPARE_GT(result, bfp_NMAX_BFP16()) then do
if OE=O then do
if x=ObOO then result \& sign ? bf p_NEGATE(bfp_INFINITY()): bf p_INFINITY()
if x=0b01 then result \& sign ? bf p_NEGATE(bf[_NMAX_BFP16()): bf __NMAX_BFP16()
if x=0b10 then result \& sign ? bf p_NEGATE(bfp_NMAX_BFP16()): bf p_l IFINITY()
if x=0bll then result \& sign ? bfp_NEGATE(bfp_INFINITY()) : bfp_NMAX_BFP16()
0x_flag \leftarrowObl
xxflag}\leftarrow0b
inc`flag\&obu
return(result)
end
else do
result.exponent \leftarrowresult.exponent - 24
ox_flag}
end
end
return(result)

```
bf p_ROUND_TO_BFP16_CEIL(x)
\(x\) is a normalized floating-point value represented in the working format, having unbounded exponent range and significand precision.

Return the smallest floating-point number having unbounded exponent range but half-precision significand precision that is greater or equal in value to \(x\).

If the magnitude of the value returned is greater than \(x, i n c_{-} f l a g\) is set to 1 .
If the value returned is not equal to \(x, x x_{\_} f \mid a g\) is set to 1 .
bf \(P_{-}\)ROUND_TO_BFP16_FLOOR( \(x\) )
\(x\) is a normalized floating-point value represented in the working format, having unbounded exponent range and significand precision.

Return the largest floating-point number having unbounded exponent range but half-precision significand precision that is lesser or equal in value to \(x\).

If the magnitude of the value returned is greater than \(x, \operatorname{inc} f l a g\) is set to 1 .
If the value returned is not equal to \(x, x_{x_{-}} f \mid a g\) is set to 1 .
bf p_ROUND_TO_BFP16_NEAR_EVEN(x)
\(x\) is a normalized floating-point value represented in the working format, having unbounded exponent range and significand precision.

Return the floating-point number having unbounded exponent range but half-precision significand precision that is nearest in value to \(x\) (in case of a tie, the floating-point number having unbounded exponent range but half-precision significand precision with the least-significant bit equal to 0 is used).

If the magnitude of the value returned is greater than \(x, \operatorname{inc} f l a g\) is set to 1 .
If the value returned is not equal to \(x, x_{x_{-}} f \mid a g\) is set to 1 .
bf \(p_{-}\)ROUND_TO_BFP16_TRUNC( \(x\) )
\(x\) is a normalized floating-point value represented in the working format, having unbounded exponent range and significand precision.

Return the largest floating-point number having unbounded exponent range but half-precision significand precision that is lesser or equal in value to \(x\) if \(x>0\), or the smallest floating-point number having unbounded exponent range but half0-precision significand precision that is greater or equal in value to \(x\) if \(x<0\).

If the magnitude of the value returned is greater than \(x, i n c_{-} f \mid a g\) is set to 1 .
If the value returned is not equal to \(x, x x_{-} f \mid a g\) is set to 1 .
bf p_ROUND_TO_I NTEGER(rmode, x)
\(x\) is a binary floating-point value that is represented in the working floating-point format and has unbounded exponent range and significand precision.

If \(x\) is an \(S N a N, v x s n a n \_f a g\) is set to 1 .
If x is a QNaN , return x .
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if \(x\) is an Infinity, return \(x\).
Otherwise, do the following.
If \(\mathrm{rmode}=0 b 000\) (Round to Nearest Even),
return the double-precision floating-point integer value that is nearest in value to \(x\) (in case of a tie, the double-precision floating-point integer value with the least-significant bit equal to 0 is used).

If \(\mathrm{rmode}=0 \mathrm{ODOL}\) (Round towards Zero),
return the largest double-precision floating-point integer value that is lesser or equal in value to \(x\) if \(x>0\), or the smallest double-precision floating-point integer value that is greater or equal in value to \(x\) if \(x<0\).

If \(\mathrm{rmode}=0 \mathrm{bOlO}\) (Round towards +Infinity),
return the smallest double-precision floating-point integer value that is greater or equal in value to \(x\).
If \(\mathrm{mode}=0 \mathrm{b011}\) (Round towards • Infinity),
return the largest double-precision floating-point integer value that is lesser or equal in value to \(x\).
If \(\mathrm{rmode}=0 \mathrm{bl} 100\) (Round to Nearest Away),
return the double-precision floating-point integer value that is nearest in value to \(x\) (in case of a tie, the double-precision floating-point integer value that is furthest away from 0 is used).
inc_flag is set to 1 if the magnitude of the value returned is greater than \(x\).
\(x x_{\mathrm{f}} \bar{f} \mathrm{ag}\) is set to 1 if the value returned is not equal to \(x\).
bf p_ROUND_ODD(p, x)
\(x\) is a binary floating-point value that is represented in the working floating-point format and has unbounded exponent range and significand precision. x must be rounded as presented, without prenormalization.
\(p\) is an integer value specifying the precision (i.e., number of bits) the significand is rounded to.
Return \(x\) with bit \(p .1\) of the significand set to 1 if any of the bits to the right of bit \(p .1\) of the significand of \(x\) are equal to 1 , and all bits to the right of bit \(p \cdot 1\) of the significand of the value returned are set to 0 . Otherwise return \(x\) with all bits to the right of bit p .1 of the significand set to 0 .
inc_flag is set to 1 if the magnitude of the value returned is greater than \(x\).
\(x x_{-} f \mid a g\) is set to 1 if the value returned is not equal to \(x\).
bf p_ROUND_NEAR_EVEN( \(p, x)\)
\(x\) is a binary floating-point value that is represented in the working floating-point format and has unbounded exponent range and significand precision. x must be rounded as presented, without prenormalization.
\(p\) is an integer value specifying the precision (i.e., number of bits) the significand is rounded to.
Return the floating-point number having unbounded exponent range and a significand with a width of \(p\) bits that is nearest in value to \(x\) (in case of a tie, the floating-point number having unbounded exponent range and a \(p\)-bit significand with the least-significant bit equal to 0 is used).
inc_flag is set to 1 if the magnitude of the value returned is greater than \(x\).
\(x x_{-} \bar{f} \mid \mathrm{ag}\) is set to 1 if the value returned is not equal to \(x\).
bf p_ROUND_TRUNC( \(p, x)\)
\(x\) is a binary floating-point value that is represented in the working floating-point format and has unbounded exponent range and significand precision. x must be rounded as presented, without prenormalization.
\(p\) is an integer value specifying the precision (i.e., number of bits) the significand is rounded to.
Return the largest floating-point number having unbounded exponent range and a significand with a width of \(p\) bits that is lesser or equal in value to \(x\) if \(x>0\), or the smallest floating-point number having unbounded exponent range but double-precision significand precision that is greater or equal in value to \(x\) if \(x<0\).
inc_flag is set to 1 if the magnitude of the value returned is greater than \(x\).
\(x x_{-} \bar{f} \operatorname{lag}\) is set to 1 if the value returned is not equal to \(x\).
bf p_ROUND_TO_BFP128(ro, rmode, x)
\(x\) is a normalized binary floating-point value that is represented in the working floating-point format and has unbounded exponent range and significand precision.
\(r 0\) is a 1-bit unsigned integer and \(r\) mode is a 2-bit unsigned integer, together specifying one of five rounding modes to be used in rounding \(z\).
\begin{tabular}{lll}
\(r 0=0\) & \(r m o d e=0 b 00\) & Round to Nearest Even \\
\(r 0=0\) & rmode \(=0 b 01\) & Round towards Zero \\
\(r 0=0\) & \(r m o d e=0 b 10\) & Round towards + Infinity \\
\(r 0=0\) & \(r m o d e=0 b 11\) & Round towards \(\operatorname{Infinity~}\) \\
\(r 0=1\) & & Round to Odd
\end{tabular}

Return the value x rounded to quad-precision under control of the specified rounding mode.
```

if x.class.QNaN then return x
if x.class.Infinity then return x
if x.class.Zero then return x
if bfp_ABSOLUTE(x)<bfp_NMIN_BFP128 then do
if FPSCR.UE=0 then do
x < bfp_DENORM(-16382,x)
if ro=0 \& rmode=0boo thenr r bfp_ROUND_NEAR_EVEN(113,x)
if ro=0 \& rmode=0bO1 then r \& bf P_ROUND_TRUNC(113,x)
if ro=0 \& rmode=0b10 then r \& bf p_ROUND_CEIL(113,x)
if ro=0 \& rmode=0b11 then r \& bf p_ROUND_FLOOR(113,x)
if ro=1 then r \& bf p_ROUND_OdD(113,x)
ux_flag\leftarrowxx_flag
return(r)
end
else do
x.exponent }\leftarrowx\mathrm{ .exponent + 24576
ux_flag}\leftarrow
end
end
if ro=0 \& rmode=0boo then r \& bfp round NEAR EVEN(113,x)
if ro=0 \& rmode=0b01 then r \& bf p_ROUND_TRUNC(113,x)
if ro=0 \& rmode=Ob10 then r \& bf p_ROUND_CEIL(113, x)
if r }0=0\mathrm{ \& rmode=0b11 then r \& bf p_ROUND_FLOOR(113,x)
if ro=1 then r \& bf p_ROUND_ODD(113,x)
if bfp_ABSOLUTE(r) >bfp_NMAX_BFP128 then do
if FPSCR.OE=O then do
if ro=0 \& rmode=ObOO thenr \& x.sign? bfp_INFINITY : bfp_INFINITY
if ro=0 \& rmode=0b01 then r \& , sign? bfp_NMAX_BPP128: bfp_NMAX_BFP128
if ro=0 \& rmode=Ob10 then r \& < sign? bfp_NMAX_bFP128: bfp_l NFINITY
if ro=0 \& rmode=0b11 then r \&x.sign? bf _INFINITY: bf _NMAX_BFP128
if ro=1 then r \& <.sign? bfp_NMAX_BFP128: bfp_NMAX_BFP128
r.sign \leftarrowx.sign
ox_flag \& Obl
xx_flag \leftarrowObl
inc_flag\leftarrowObU
return(r)
end
else do
r.exponent \&r.exponent - 24576
Ox_flag}\leftarrow
end
end
return(r)

```
bf \(p_{-}\)ROUND_TO_BFP80(rmode, \(x\) )
\(x\) is a normalized binary floating-point value that is represented in the working floating-point format and has unbounded exponent range and significand precision.
\(r\) mode is a 2-bit unsigned integer, together specifying one of four rounding modes to be used in rounding \(x\).
\[
\begin{array}{ll}
\mathrm{rmode}=0 \mathrm{boO} & \text { Round to Nearest Even } \\
\mathrm{rmode}=0 \mathrm{bO1} & \text { Round towards Zero } \\
\mathrm{rmode}=0 \mathrm{~b} 10 & \text { Round towards + Infinity } \\
\mathrm{rmode}=0 \mathrm{~b} 11 & \text { Round towards } \text { - Infinity }
\end{array}
\]

Return the value \(x\) rounded to double-extended-precision under control of the specified rounding mode.
```

if x.class.QNaN then return x
if x.class.Infinity then return x
if x.class.Zero then return x
if bfp_ABSOLUTE (x)<bfp_NMIN_BFP80 then do
if FPSCR,UE=O then do
x \&bfp_DENORM(-16382,x)
if rmode=0b00 then r \& bf p_ROUND_NEAR_EVEN(64,x)
if rmode=ObO1 then r \& bf p_ROUND_TRUNC( 64, x)
if rmode=Ob10 then r \& bf__ROUND_CEIL(64,x)
if rmode=Ob11 then r \& bfp_ROUND_FLOOR(64,x)
uxffag < xx_flag
return(r)
end
else do
x.exponent }\leftarrowx\mathrm{ .exponent + 24576
ux_flag \& 1
end
end
if rmode=0b00 then r \& bf p_ROUND_NEAR_EVEN(64,x)
if rmode=ObO1 then r \& bf __ROUND_TRUNC( 64, x)
if rmode=0b10 then r \& bf p_ROUND_CEIL(64,x)
if rmode=0bl1 then r \&bfp_ROUND_FLOOR(64,x)
if bfp_ABSOLUTE(r)>bfp_NMAX_BFP8O then do
if FPSCR,OE=O then do
if rmode=ObOO then r \&x.sign? bf P_INFINITY: bf P_ INFIN|TY
if rmode=0b01 then r \& < sign? bf p_NmAX_BFP80: bf p_NmAX_BFP80
if rmode=Ob10 then r \& x.sign? bf p_NMAX_BFP80: bf p_l NFINITY
if rmode=Obll then r \&x.sign? bf p_INFINITY : bf p_NMAX_BFP8O
r.sign }\leftarrowx.sig
Ox_flag \leftarrowObl
xx_flag \leftarrowObl
inc_flag}\leftarrow0b
return(r)
end
else do
r.exponent \&r.exponent - 24576
Ox_flag \& 1
end
end
return(r)

```
bf p_ROUND_TO_BFP64(ro, rmode, x)
\(x\) is a normalized binary floating-point value that is represented in the working floating-point format and has unbounded exponent range and significand precision.
\(r 0\) is a 1-bit unsigned integer and \(r\) mode is a 2-bit unsigned integer, together specifying one of five rounding modes to be used in rounding \(z\).
\begin{tabular}{lll}
\(r 0=0\) & \(r m o d e=0 b 00\) & Round to Nearest Even \\
\(r 0=0\) & rmode \(=0 b 01\) & Round towards Zero \\
\(r 0=0\) & \(r m o d e=0 b 10\) & Round towards + Infinity \\
\(r 0=0\) & \(r m o d e=0 b 11\) & Round towards \(\operatorname{Infinity~}\) \\
\(r 0=1\) & & Round to Odd
\end{tabular}

Return the value \(x\) rounded to double-precision under control of the specified rounding mode.
```

if x.class.QNaN then return x
if x.class.Infinity then return x
if x.class.Zero then return x
if bfp_ABSOLUTE(x)<bfp_NMIN_BFP64 then do
if FPSCR.UE=0 then do
x}\leftarrowbfp_DENORM(-1022,x
if ro=0 \& rmode=0boo thenr r bf p_ROUND_NEAR_Even( 53,x)
if ro=0 \& rmode=0bO1 then r <bf P_ROUND_TRUNC(53,x)
if ro=0 \& rmode=0b10 then r \& bf p_ROUND_CEIL(53,x)
if ro=0 \& rmode=Ob11 then r \&bf p_ROUND_FLOOR(53,x)
if r }0=1\quad\mathrm{ then r \& bf _ _round_ODD(53,x)
uxflag}\leftarrowx\mp@subsup{x}{_}{\primeflag
return(r)
end
else do
x. exponent \leftarrowx. exponent + 1536
ux_flag}\leftarrow
end
end
if ro=0 \& rmode=0boo thenr \& bf p ROUND NEAR EVEN( 53, x)
if ro=0 \& rmode=0bO1 then r \&bfp_ROUND TRUNC(53,x)
if ro=0 \& rmode=Ob10 then r \& bf p_ROUND_CEIL(53,x)
if ro=0 \& rmode=Ob11 then r \&bf P_ROUND_FLOOR(53,x)
if ro=1 then r \&bf_ROUND_ODD(53,x)
if bfp_ABSOLUTE(r)>bfp_NMAX_BFP64 then do
if FPSCR.OE=O then do
if ro=0 \& rmode=ObOO then r \& x.sign? bfp_ INFINITY : bfp_I NFINI TY
if ro=0 \& rmode=0b01 then r \&x.sign? bf _NMAX_BFP64: bfp_NMAX_BFP64

```

```

            if ro=0 & rmode=0b11 then r &x.sign? bfp_INFINITY : bfp_NMAX_BFP64
            if ro=1 then r &x.sign? bfp_NMAX_BFP64: bf p_NMAX_BFP64
            r.sign }\leftarrowx.sig
            Ox_flag \leftarrowObl
            xx_flag \leftarrowObl
            inc_flag\leftarrowObU
            return(r)
        end
        else do
            r.exponent &r.exponent · 1536
            Ox_flag}\leftarrow
        end
    end
return(r)

```
bf P_SQUARE_ROOT(x)
\(x\) is a binary floating-point value that is represented in the working floating-point format and has unbounded exponent range and significand precision.

If \(x\) is an \(S N a N, v \times s n_{n-} f l a g\) is set to 1 .
Otherwise, if x is negative and non-zero, vxsqrt fl ag is set to 1 .
If x is a QNaN , return x .
Otherwise, if \(x\) is an \(S N a N\), return \(x\) represented as a QNaN .
Otherwise, if \(x\) is - Zero, return - Zero.
Otherwise, if \(x\) is negative, return the standard QNaN .
Otherwise, return the normalized square root of \(x\), having unbounded range and precision.

\section*{ClassDP(x,y)}

Return a 5-bit characterization of the double-precision floating-point number \(x\).
```

Ob10001 = Quiet NaN
0b01001 = -Infinity
0b01000 = -Normalized Number
0b11000 $=-$-Denormalized Number
Ob10010 = -Zero
0b00010 = +Zero
0b10100 = +Denormalized Number
0b00100 $=+$ Normalized Number
0b00101 $=+$ Infinity

```

\section*{ClassSP(x,y)}

Return a 5-bit characterization of the single-precision floating-point number \(x\).
```

Ob10001 = Quiet NaN
0b01001 = -Infinity
0b01000 = -Normalized Number
0b11000 = -Denormalized Number
Ob10010 = -Zero
0b00010 = +Zero
0b10100 = +Denormalized Number
0b00100 = +Normalized Number
0b00101 = +Infinity

```

\section*{CompareEQDP( \(\mathbf{x , y}\) )}
\(x\) and \(y\) are double-precision floating-point values.
If \(x\) or \(y\) is a \(N a N\), return 0 .
Otherwise, if \(x\) is equal to \(y\), return 1 .
Otherwise, return 0.

\section*{CompareEQSP( \(\mathbf{x}, \mathrm{y}\) )}
\(x\) and \(y\) are single-precision floating-point values.
If \(x\) or \(y\) is a NaN , return 0 ,
Otherwise, if \(x\) is equal to \(y\), return 1 .
Otherwise, return 0.

\section*{CompareGTDP( \(\mathbf{x , y}\) )}
\(x\) and \(y\) are double-precision floating-point values.
If \(x\) or \(y\) is a \(N a N\), return 0 ,
Otherwise, if x is greater than y , return 1 .
Otherwise, return 0 .

\section*{CompareGTSP( \(\mathbf{x}, \mathrm{y}\) )}
\(x\) and \(y\) are single-precision floating-point values.
If x or y is a NaN , return 0 .
Otherwise, if x is greater than y , return 1 .
Otherwise, return 0.

\section*{CompareLTDP( \(\mathrm{x}, \mathrm{y}\) )}
\(x\) and \(y\) are double-precision floating-point values.
If \(x\) or \(y\) is a NaN, return 0 .
Otherwise, if x is less than y , return 1 .
Otherwise, return 0.

\section*{CompareLTSP( \(\mathbf{x}, \mathrm{y}\) )}
\(x\) and \(y\) are single-precision floating-point values.
If \(x\) or \(y\) is a NaN, return 0 .
Otherwise, if x is less than y , return 1 .
Otherwise, return 0.

\section*{ConvertDPtoSD(x)}
\(x\) is a floating-point value in double-precision format.
If x is a NaN ,
vxcui_flag is set to 1 ,
vxsnan_flag is set to 1 if \(x\) is an SNaN, and
return \(0 \times 8000 \_0000 \_0000 \_0000\),
Otherwise, do the following.
Let \(r\) nd be the value \(x\) truncated to an integral value.
If \(r n d\) is greater than \(2^{63} \cdot 1\),
vxcui_flag is set to 1 ,
return \(0 \times 7 F F F\) FFFF_FFFF_FFF.
Otherwise, if \(\mathrm{r} n \mathrm{~d}\) is less than \(-2^{63}\),
vxcui_flag is set to 1 ,
return \(0 \times 8000 \_0000 \_0000 \_0000\).
Otherwise,
\(x x_{-} f \mid a g\) is set to 1 if \(r n d\) is inexact.
return \(r\) nd in 64-bit signed integer format.

\section*{ConvertDPtoSP(x)}
x is a floating-point value in double-precision format.
If \(x\) is an SNaN, vxsnan_flag is set to 1 .
If x is a SNaN , returns x , converted to a QNaN , in single-precision floating-point format.
Otherwise, if x is a QNaN, an Infinity, or a Zero, returns x in single-precision floating-point format.
Otherwise, returns \(x\), rounded to single-precision using the rounding mode specified in RN, in single-precision floating-point format.
\(0 \times \_f a g\) is set to 1 if rounding \(\times\) resulted in an Overflow exception.
\(u x_{-} f \mid a g\) is set to 1 if rounding \(x\) resulted in an Underflow exception.
\(x x_{-} f \mid a g\) is set to 1 if rounding \(x\) returns an inexact result.
inc_flag is set to 1 if the significand of the result was incremented during rounding.

\section*{ConvertDPtoSP_NS(x)}
\(x\) is a single-precision floating-point value represented in double-precision format.
Returns x in single-precision format.
```

sign }\leftarrowx.bit[0
exponent \leftarrowx.bit[1:11]
fraction \leftarrow 0b1 || x.bit[12:63] // implicit bit set to 1 (for now)
if (exponent == 0) \& (fraction.bit[1:52] != 0) then do // DP Denormal operand
exponent \leftarrow 0b000_0000_0001 // exponent override to DP Emin = 1
fraction.bit[0] \leftarrow 0b0 // implicit bit override to 0
end
if (exponent < 897) \&\& (fraction != 0) then do // SP tiny operand
fraction \leftarrow fraction >> ui (897 - exponent) // denormalize until exponent = SP Emin
exponent \leftarrow 0b011_1000_0000 // exponent override to SP Emin-1 = 896
end
return(sign || exponent.bit[0] | exponent.bit[4:10] | fraction.bit[1:23])

```

\section*{Programming Note}

If \(x\) is not representable in single-precision, some exponent and/or significand bits will be discarded, likely producing undesirable results. The low-order 29 bits of the significand of \(x\) are discarded, more if the unbiased exponent of \(x\) is less than - 126 (i.e., denormal). Finite values of \(x\) having an unbiased exponent less than \(\cdot 150\) will return a result of Zero. Finite values of \(x\) having an unbiased exponent greater than +127 will result in discarding significant bits of the exponent. SNaN inputs having no significant bits in the upper 23 bits of the signifcand will return Infinity as the result. No status is set for any of these cases.

\section*{ConvertDPtoSW(x)}
\(x\) is a floating-point value in double-precision format.
If \(x\) is a NaN ,
vxcri_flag is set to 1 ,
vxsnan_flag is set to 1 if \(x\) is an SNaN, and
return \(0 \times 8000 \_0000\),
Otherwise, do the following.
Let \(r n d\) be the value \(x\) truncated to an integral value.
If \(r n d\) is greater than \(2^{31} .1\), vxcui_flag is set to 1 , return \(0 \times\) PFFF_FFFF.

Otherwise, if r nd is less than \(\cdot 2^{31}\),
vxcui_flag is set to 1 , return \(0 \times 80000000\).

Otherwise, \(x x_{-} f l a g\) is set to 1 if \(r n d\) is inexact. return r nd in 32-bit signed integer format.

\section*{Version 3.0}

\section*{ConvertDPtoUD(x)}
\(x\) is a floating-point value in double-precision format.
If x is a NaN ,
vxcui_flag is set to 1,
vxsnan_flag is set to 1 if \(x\) is an SNaN , and return \(0 \times 8000 \_0000 \_0000 \_0000\),

Otherwise, do the following.
Let \(r n d\) be the value \(x\) truncated to an integral value.
If \(r n d\) is greater than \(2^{64.1}\),
vxcui_flag is set to 1 ,
return \(0 \times\) FFFF_FFFFFFFF_FFF.
Otherwise, if \(r\) nd is less than 0 ,
vxcui_flag is set to 1 ,
return \(0 \times 0000 \_0000 \_0000 \_0000\).
Otherwise,
\(x x_{-} f \mid a g\) is set to 1 if \(r n d\) is inexact.
return \(r\) nd in 64-bit unsigned integer format.

\section*{ConvertDPtoUW(x)}
\(x\) is a floating-point value in double-precision format.
If \(x\) is a NaN ,
vxcui_flag is set to 1 ,
vxsnan_flag is set to 1 if \(x\) is an SNaN, and
return \(0 \times 0000 \_0000\),
Otherwise, do the following.
Let r nd be the value x truncated to an integral value.
If \(r n d\) is greater than \(2^{32} \cdot 1\),
vxcui_flag is set to 1 ,
return \(0 \times F F F F\) FFFF.
Otherwise, if rnd is less than 0 ,
vxcui_flag is set to 1 ,
return \(0 \times 0000 \_0000\).
Otherwise,
\(x x_{\_} f \mid a g\) is set to 1 if \(r n d\) is inexact.
return \(r\) nd in 32-bit unsigned integer format.

\section*{ConvertFPtoDP(x)}

Return the floating-point value x in DP format.

\section*{ConvertFPtoSP(x)}

Return the floating-point value x in single-precision format.

\section*{ConvertSDtoFP(x)}
\(x\) is a 64-bit signed integer value.
Return the value x converted to floating-point format having unbounded significand precision.

\section*{ConvertSPtoDP_NS(x)}
\(x\) is a single-precision floating-point value.
Returns x in double-precision format.
```

sign }\leftarrowx.bit[0
exponent \leftarrow (x.bit[1] || \negx.bit[1] | \negx.bit[1] || \negx.bit[1] | x.bit[2:8])
fraction \leftarrow 0b0 || x.bit[9:31] || 0b0_0000_0000_0000_0000_0000_0000_0000
if (x.bit[1:8] == 255) then do // Infinity or NaN operand
exponent \leftarrow2047 // override exponent to DP Emax+1
end
else if (x.bit[1:8] == 0) \&\& (fraction == 0) then do // SP Zero operand
exponent \leftarrow 0
end
else if (x.bit[1:8] == 0) \&\& (fraction != 0) then do // SP Denormal operand
exponent \leftarrow897
do while (fraction.bit[0] == 0)
// override exponent to SP Emin
fraction \leftarrow fraction << 1
exponent }\leftarrow\mathrm{ exponent - 1
end
end
return(sign || exponent || fraction.bit[1:52])

```

\section*{Version 3.0}

\section*{ConvertSP64toSP(x)}
x is a single-precision floating-point value in double-precision format.
Returns the value \(x\) in single-precision format. x must be representable in single-precision, or else result returned is undefined. \(x\) may require denormalization. No rounding is performed. If \(x\) is a SNaN , it is converted to a sin-gle-precision SNaN having the same payload as x .
```

sign \leftarrow x.bit[0]
exp \leftarrow x.bit[1:11] - 1023
frac }\leftarrow x.bit[12:63
if (exp = -1023) \& (frac = 0) \& (sign=0) then return(0x0000_0000) // +Zero
else if (exp = -1023) \& (frac = 0) \& (sign=1) then return(0x8000_0000) // -Zero
else if (exp = -1023) \& (frac != 0) then return(0xUUUU_UUUU) // DP denorm
else if (exp < -126) then do // denormalization required
msb = 1
do while (exp < -126) // denormalize operand until exp=Emin
frac.bit[1:51] \leftarrow frac.bit[0:50]
frac.bit[0] \leftarrow msb
msb }\leftarrow
exp }\leftarrow\operatorname{exp}+
end
if (frac = 0) then return(0xUUUU_UUUU) // value not representable in SP format
else do // return denormal SP
result.bit[0] }\leftarrow\mathrm{ sign
result.bit[1:8] \leftarrow 0
result.bit[9:31] \leftarrow frac.bit[0:22]
return(result)
end
end
else if (exp = +1024) \& (frac = 0) \& (sign=0) then return(0x7F80_0000) // +Infinity
else if (exp = +1024) \& (frac = 0) \& (sign=1) then return(0xFF80_0000) // -Infinity
else if (exp = +1024) \& (frac != 0) then do // QNaN or SNaN
result.bit[0] \leftarrow sign
result.bit[1:8] }\leftarrow25
result.bit[9:31] \leftarrow frac.bit[0:22]
return(result)
end
else if (exp < +1024) \& (exp > +126) then return(0xUUUU_UUUU) // overflow
else do // normal value
result.bit[0] \leftarrow sign
result.bit[1:8] \leftarrow exp.bit[4:11] + 127
result.bit[9:31] \leftarrow frac.bit[0:22]
return(result)
end

```

\section*{ConvertSPtoDP(x)}
\(x\) is a single-precision floating-point value.
If \(x\) is an \(\mathrm{SNaN}, \mathrm{vxsnan} \_f l a g\) is set to 1 .
If \(x\) is an SNaN , return x represented as a QNaN in double-precision floating-point format. Otherwise, if \(x\) is an \(Q N a N\), return \(x\) in double-precision floating-point format.
Otherwise, return the value \(x\) in double-precision floating-point format.

\section*{ConvertSPtoSD(x)}
x is a floating-point value in single-precision format.
If x is a NaN ,
vxcvi_flag is set to 1 , and
vxsnan_flag is set to 1 if \(x\) is an SNaN return \(0 \times 8000 \_0000 \_0000 \_0000\) and

Otherwise, do the following.
Let \(r\) nd be the value \(x\) truncated to an integral value.
If \(r n d\) is greater than \(2^{63} \cdot 1\), vxcri_flag is set to 1 , and return \(0 \times 7\) FFF_FFFFFFFF_FFF.

Otherwise, if r nd is less than \(\cdot 2^{63}\), vxcui_flag is set to 1 , and return \(0 \times 8000 \_0000 \_0000 \_0000\).

Otherwise,
\(x x_{-} f \mid a g\) is set to 1 if \(r n d\) is inexact, and return \(r\) nd in 64-bit signed integer format.

\section*{ConvertSPtoSP64(x)}
x is a floating-point value in single-precision format.
Returns the value \(x\) in double-precision format. If \(x\) is a SNaN, it is converted to a double-precision \(\operatorname{SNaN}\) having the same payload as \(x\).
```

sign }\leftarrowx.bit[0
exp \leftarrowx.bit[1:8] - 127
frac}\leftarrowx.bit[9:31
if (exp = -127) \& (frac != 0) then do || Normalize the Denormal value
msb}\leftarrow\textrm{frac.bit[0]
frac \leftarrowfrac << l
do while (msb = 0)
msb}\leftarrow\textrm{frac.bit[0]
frac}\leftarrowfrac<< l
exp \leftarrowexp - 1
end
end
else if (exp = - 127) \& (frac = 0) then exp \leftarrow-1023 || Zero value
else if (exp = +128) then exp \leftarrow +1024 || |nfinity, NaN
result.bit[0]}\leftarrow\mathrm{ sign
result.bit[1:11]}\leftarrow\operatorname{exp + 1023
result.bit[12:34]}\leftarrow\textrm{frac
result.bit[35:63]}\leftarrow
return(result)

```

\section*{Version 3.0}

\section*{ConvertSPtoSW(x)}
\(x\) is a floating-point value in single-precision format.
If x is a NaN ,
vxcui_flag is set to 1 ,
vxsnan_flag is set to 1 if \(x\) is an SNaN , and return \(0 \times 80000000\).

Otherwise, do the following.
Let r nd be the value x truncated to an integral value.
If \(r n d\) is greater than \(2^{31} \cdot 1\), vxcvi_flag is set to 1 , and return \(0 \times 7 F F F\) FFFF.

Otherwise, if rnd is less than \(\cdot 2^{31}\), vxcvi_flag is set to 1 , and return \(0 \times 8000 \_0000\).

Otherwise,
\(x x_{-} f l a g\) is set to 1 if \(r n d\) is inexact, and return \(r\) nd in 32-bit signed integer format.

\section*{ConvertSPtoUD(x)}
x is a floating-point value in single-precision format.
If x is a NaN ,
vxcvi_flag is set to 1 , and
vxsnan_flag is set to 1 if \(x\) is an SNaN
return \(0 \times 0000 \_0000 \_0000 \_0000\),
Otherwise, do the following.
Let r nd be the value x truncated to an integral value.
If \(r n d\) is greater than \(2^{64} \cdot 1\), vxcvi_flag is set to 1 , and return \(0 \times F F F F\) FFFF_FFFF_FFF.

Otherwise, if rnd is less than 0 , vxcui_flag is set to 1 , and return \(0 \times 0000 \_0000 \_0000 \_0000\).

Otherwise,
\(x x_{\_} f l a g\) is set to 1 if \(r n d\) is inexact, and return \(r n d\) in 64-bit unsigned integer format.

\section*{ConvertSPtoUW(x)}
\(x\) is a floating-point value in single-precision format.
If x is a NaN , vxcui_flag is set to 1 , vxsnan_flag is set to 1 if \(x\) is an SNaN , and return \(0 \times 0000 \_0000\).

Otherwise, do the following.
Let \(r\) nd be the value \(x\) truncated to an integral value.
If \(r n d\) is greater than \(2^{32} \cdot 1\), vxcui_flag is set to 1 , and return OXFFFF_FFFF.

Otherwise, if \(r n d\) is less than 0 , vxcui_flag is set to 1 , and return \(0 \times 0000 \_0000\).

Otherwise,
\(x x_{-} f \mid a g\) is set to 1 if \(r n d\) is inexact, and return \(r\) nd in 32-bit unsigned integer format.

\section*{ConvertSWtoFP(x)}
\(x\) is a 32-bit signed integer value.
Return the value x converted to floating-point format having unbounded significand precision.

\section*{ConvertUDtoFP(x)}
\(x\) is a 64-bit unsigned integer value.
Return the value x converted to floating-point format having unbounded significand precision.

\section*{ConvertUWtoFP(x)}
x is a 32-bit unsigned integer value.
Return the value x converted to floating-point format having unbounded significand precision.

\section*{DivideDP( \(x, y\) )}
\(x\) and \(y\) are double-precision floating-point values.
If \(x\) or \(y\) is an \(\mathrm{SNaN}, \mathrm{vxsnan} \mathrm{\_flag} \mathrm{is} \mathrm{set} \mathrm{to} 1\).
If x is a Zero and y is a Zero, vxzdz _flag is set to 1 .
If \(x\) is a finite, nonzero value and \(y\) is a Zero, \(z x \_f l a g\) is set to 1 .
If \(x\) is an Infinity and \(y\) is an Infinity, vxidi_flag is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an \(S N a N\), return \(x\) represented as a QNaN.
Otherwise, if y is a QNaN , return y .
Otherwise, if \(y\) is an SNaN , return y represented as a QNaN.
Otherwise, if \(x\) is a Zero and \(y\) is a Zero, return the standard QNaN .
Otherwise, if x is a finite, nonzero value and y is a Zero with the same sign as x , return + Infinity.
Otherwise, if x is a finite, nonzero value and y is a Zero with the opposite sign as x , return -Infinity.
Otherwise, if \(x\) is an Infinity and \(y\) is an Infinity, return the standard QNaN.
Otherwise, return the normalized quotient of \(x\) divided by \(y\), having unbounded range and precision.

\section*{DivideSP(x,y)}
\(x\) and \(y\) are single-precision floating-point values.
If \(x\) or \(y\) is an \(S N a N, v x s n a n \_f l a g\) is set to 1 .

If \(x\) is a finite, nonzero value and \(y\) is a Zero, \(z x \_f l a g\) is set to 1 .
If x is an Infinity and y is an Infinity, vxidi_flag is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if x is an SNaN , return x represented as a QNaN .
Otherwise, if \(y\) is a QNaN , return y .
Otherwise, if y is an SNaN , return y represented as a QNaN .
Otherwise, if x is a Zero and y is a Zero, return the standard QNaN .
Otherwise, if x is a finite, nonzero value and y is a Zero with the same sign as x , return + Infinity.
Otherwise, if \(x\) is a finite, nonzero value and \(y\) is a Zero with the opposite sign as \(x\), return -Infinity.
Otherwise, if \(x\) is an Infinity and \(y\) is an Infinity, return the standard QNaN.
Otherwise, return the normalized quotient of \(x\) divided by \(y\), having unbounded range and precision.

\section*{DenormDP(x)}
\(x\) is a floating-point value having unbounded range and precision.
Return the value \(x\) with its significand shifted right by a number of bits equal to the difference of the -1022 and the unbiased exponent of \(x\), and its unbiased exponent set to -1022 .

\section*{DenormSP(x)}
\(x\) is a floating-point value having unbounded range and precision.
Return the value \(x\) with its significand shifted right by a number of bits equal to the difference of the -126 and the unbiased exponent of \(x\), and its unbiased exponent set to -126 .

\section*{EXTZ32(x)}

Result of extending the \(b\)-bit value \(x\) on the left with \(32 \cdot b\) zeros, forming a 32 -bit value.
```

b}\leftarrow\operatorname{LENGTH(x)
result.bit[0:31-b]}\leftarrow
result.bit[32-b:31]}\leftarrow

```

\section*{EXTZ64(x)}

Result of extending the \(b\)-bit value \(x\) on the left with \(64 \cdot b\) zeros, forming a 64 -bit value.
```

b}\leftarrow\operatorname{LENGTH(x)
result.bit[0:63-b]}\leftarrow
result.bit[64-b:63]}\leftarrow

```

EXTZ128(x)
Result of extending the \(b\)-bit value \(x\) on the left with \(128 \cdot b\) zeros, forming a 128-bit value.
```

b}\leftarrow\operatorname{LENGTH(x)
result.bit[0:127-b]}\leftarrow
result.bit[128.b:127]}\leftarrow

```

\section*{fprf_CLASS_BFP16(x)}
\(x\) is a floating-point value represented in half-precision format.
Return the 5 -bit code that specifies the sign and class of \(x\).
Return \(0 b 10001\) if x is a Quiet NaN .
Return \(0 b 01001\) if \(x\) is a negative infinity.
Return \(0 b 00101\) if x is a positive infinity.
Return \(0 b 10010\) if \(x\) is a negative zero.
Return \(0 b 00010\) if x is a positive zero.
Return \(0 b 11000\) if x is a negative denormal value when represented in half-precision format.
Return 0 b10100 if x is a positive denormal value when represented in half-precision format.
Return 0601000 if x is a negative normal value when represented in half-precision format.
Return \(0 b 00100\) if x is a positive normal value when represented in half-precision format.
fprf_CLASS_BFP64(x)
\(x\) is a floating-point value represented in double-precision format.
Return the 5-bit code that specifies the sign and class of \(x\).
Return \(0 b 10001\) if x is a Quiet NaN .
Return 0b01001 if \(x\) is a negative infinity.
Return 0b00101 if x is a positive infinity.
Return \(0 b 10010\) if \(x\) is a negative zero.
Return \(0 b 00010\) if x is a positive zero.
Return \(0 b 11000\) if \(x\) is a negative denormal value when represented in double-precision format.
Return \(0 b 10100\) if \(x\) is a positive denormal value when represented in double-precision format.
Return 0601000 if \(x\) is a negative normal value when represented in double-precision format.
Return \(0 b 00100\) if x is a positive normal value when represented in double-precision format.
fprf_CLASS_BFP128(x)
\(x\) is binary floating-point value that is represented in quad-precision format.
Return the 5-bit characterization of the sign and class of \(x\).
Return \(0 b 10001\) if \(x\) is a Quiet NaN.
Return \(0 b 01001\) if \(x\) is negative and an infinity.
Return \(0 b 01000\) if \(x\) is negative and a normal number.
Return 0 b 11000 if x is negative and a denormal number.
Return 0 b 10010 if x is negative and a zero.
Return 0b00010 if \(x\) is positive and a zero.
Return \(0 b 10100\) if x is positive and a denormal number.
Return \(0 b 00100\) if x is positive and a normal number.
Return \(0 b 00101\) if x is positive and an infinity.
\(\operatorname{Is} \operatorname{Inf}(x)\)
Return 1 if x is an Infinity, otherwise return 0.

\section*{IsNaN(x)}

Return 1 if x is either an SNaN or a QNaN , otherwise return 0 .
IsNeg(x)
Return 1 if x is a negative, nonzero value, otherwise return 0 .

\section*{IsSNaN(x)}

Return 1 if x is an SNaN , otherwise return 0 .

\section*{IsZero(x)}

Return 1 if \(x\) is a Zero, otherwise return 0 .

\section*{MaximumDP( \(x, y\) )}
\(x\) and \(y\) are double-precision floating-point values.
If \(x\) or \(y\) is an \(S N a N, v x s n a n \_f l a g\) is set to 1 .
If \(x\) is a \(Q N a N\) and \(y\) is not a \(N a N\), return \(y\).
Otherwise, if \(x\) is a QNaN, return \(x\).
Otherwise, if \(x\) is an \(S N a N\), return \(x\) represented as a QNaN.
Otherwise, if y is a QNaN , return x .
Otherwise, if y is an SNaN , return y represented as a QNaN .
Otherwise, return the greater of x and y , where +0 is considered greater than -0 .

\section*{MaximumSP(x,y)}
\(x\) and \(y\) are single-precision floating-point values.
If \(x\) or \(y\) is an SNaN , vxsnan_flag is set to 1 .
If \(x\) is a \(Q N a N\) and \(y\) is not a NaN, return \(y\).
Otherwise, if \(x\) is a QNaN, return \(x\).
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if y is a QNaN , return x .
Otherwise, if y is an SNaN , return y represented as a QNaN.
Otherwise, return the greater of \(x\) and \(y\), where +0 is considered greater than -0 .

\section*{MinimumDP( \(\mathrm{x}, \mathrm{y}\) )}
\(x\) and \(y\) are double-precision floating-point values.
If \(x\) or \(y\) is an \(S N a N\), \(v x\) snan_flag is set to 1 .
If \(x\) is a \(Q N a N\) and \(y\) is not a \(N a N\), return \(y\).
Otherwise, if \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an \(S N a N\), return \(x\) represented as a QNaN.
Otherwise, if y is a QNaN , return x .
Otherwise, if y is an SNaN , return y represented as a QNaN.
Otherwise, return the lesser of \(x\) and \(y\), where -0 is considered less than +0 .

\section*{MinimumSP( \(x, y\) )}
\(x\) and \(y\) are single-precision floating-point values.
If \(x\) or \(y\) is an SNaN, vxsnan_flag is set to 1 .
If \(x\) is a \(Q N a N\) and \(y\) is not a \(N a N\), return \(y\).
Otherwise, if \(x\) is a QNaN, return \(x\).
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if \(y\) is a QNaN , return x .
Otherwise, if y is an SNaN , return y represented as a QNaN .
Otherwise, return the lesser of \(x\) and \(y\), where -0 is considered less than +0 .

\section*{MultiplyAddDP(x,y,z)}
\(x, y\) and \(z\) are double-precision floating-point values.
If \(x, y\) or \(z\) is an \(S N a N, v x\) nan_flag is set to 1 .
If x is a Zero and y , is an Infinity or x is an Infinity and y is an Zero, vximz_flag is set to 1 .
If the product of \(x\) and \(y\) is an Infinity and \(z\) is an Infinity of the opposite sign, vxisi_flag is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an \(S N a N\), return \(x\) represented as a \(Q N a N\).
Otherwise, if \(z\) is a QNaN , return \(z\).
Otherwise, if \(z\) is an SNaN , return \(z\) represented as a QNaN .
Otherwise, if y is a QNaN , return y .
Otherwise, if y is an SNaN , return y represented as a QNaN .
Otherwise, if x is a Zero and y is an Infinity or x is an Infinity and y is an Zero, return the standard QNaN .
Otherwise, if the product of \(x\) and \(y\) is an Infinity, and \(z\) is an Infinity of the opposite sign, return the standard QNaN.
Otherwise, return the normalized sum of \(z\) and the product of \(x\) and \(y\), having unbounded range and precision.

\section*{MultiplyAddSP(x,y,z)}
\(x, y\) and \(z\) are single-precision floating-point values.
If \(x, y\) or \(z\) is an \(S N a N\), vxsnan_flag is set to 1 .
If x is a Zero and y is an Infinity, or x is an Infinity and y is an Zero, vximz_flag is set to 1.
If the product of \(x\) and \(y\) is an Infinity and \(z\) is an Infinity of the opposite sign, vxisi_flag is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an \(S N a N\), return \(x\) represented as a QNaN.
Otherwise, if \(z\) is a \(Q N a N\), return \(z\).
Otherwise, if \(z\) is an \(S N a N\), return \(z\) represented as a QNaN.
Otherwise, if y is a QNaN , return y .
Otherwise, if \(y\) is an SNaN , return y represented as a QNaN.
Otherwise, if \(x\) is a Zero and \(y\) is an Infinity or \(x\) is an Infinity and \(y\) is an Zero, return the standard QNaN.
Otherwise, if the product of \(x\) and \(y\) is an Infinity, and \(z\) is an Infinity of the opposite sign, return the standard QNaN.
Otherwise, return the normalized sum of \(z\) and the product of \(x\) and \(y\), having unbounded range and precision.
MultiplyDP(x,y)
\(x\) and \(y\) are double-precision floating-point values.
If \(x\) or \(y\) is an SNaN, vxsnan_flag is set to 1 .

If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an \(S N a N\), return \(x\) represented as a QNaN.
Otherwise, if y is a QNaN , return y .
Otherwise, if \(y\) is an SNaN , return y represented as a QNaN .
Otherwise, if x is a Zero and y is as Infinity or x is a Infinity and y is an Zero, return the standard QNaN .
Otherwise, return the normalized product of \(x\) and \(y\), having unbounded range and precision.

\section*{Version 3.0}

\section*{MultiplySP(x,y)}
\(x\) and \(y\) are single-precision floating-point values.
If \(x\) or \(y\) is an \(S N a N, v x s n a n \_f l a g\) is set to 1 .
If x is a Zero and y is an Infinity, or x is an Infinity and y is an Zero, vximz_flag is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN.
Otherwise, if y is a QNaN, return y .
Otherwise, if \(y\) is an SNaN , return y represented as a QNaN.
Otherwise, if x is a Zero and y is as Infinity or x is a Infinity and y is an Zero, return the standard QNaN .
Otherwise, return the normalized product of \(x\) and \(y\), having unbounded range and precision.

\section*{NegateDP(x)}

If the double-precision floating-point value x is a NaN , return x .
Otherwise, return the double-precision floating-point value x with its sign bit complemented.

\section*{NegateSP(x)}

If the single-precision floating-point value x is a NaN , return x .
Otherwise, return the single-precision floating-point value x with its sign bit complemented.

\section*{ReciprocalEstimateDP(x)}
x is a double-precision floating-point value.
If \(x\) is an \(S N a N, v x s n a n \_f l a g\) is set to 1 .
If \(x\) is a Zero, \(z x \_f l a g\) is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN.
Otherwise, if \(x\) is a Zero, return an Infinity with the sign of \(x\).
Otherwise, if \(x\) is an Infinity, return a Zero with the sign of \(x\).
Otherwise, return an estimate of the reciprocal of \(x\) having unbounded exponent range.

\section*{ReciprocalEstimateSP(x)}
x is a single-precision floating-point value.
If \(x\) is an \(S N a N, v x s n a n \_f l a g\) is set to 1 .
If \(x\) is a Zero, \(z x \_f l a g\) is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN.
Otherwise, if \(x\) is a Zero, return an Infinity with the sign of \(x\).
Otherwise, if \(x\) is an Infinity, return a Zero with the sign of \(x\).
Otherwise, return an estimate of the reciprocal of \(x\) having unbounded exponent range.
```

ReciprocalSquareRootEstimateDP(x)
x is a double-precision floating-point value.
If }x\mathrm{ is an SNaN, vxsnan_flag is set to 1.
If x is a Zero, zx_flag is set to 1.
If }x\mathrm{ is a negative, nonzero number, vxsqrt_flag is set to 1.
If }\textrm{x}\mathrm{ is a QNaN, return x.
Otherwise, if x is an SNaN, return x represented as a QNaN.
Otherwise, if x is a negative, nonzero value, return the default QNaN.
Otherwise, return an estimate of the reciprocal of the square root of x having unbounded exponent range.
ReciprocalSquareRootEstimateSP(x)
x is a single-precision floating-point value.
If }x\mathrm{ is an SNaN, vxsnan_flag is set to 1.
If x is a Zero, zx_flag is set to 1.
If }x\mathrm{ is a negative, nonzero number, vxsqrt_flag is set to 1.
If }\textrm{x}\mathrm{ is a QNaN, return x.
Otherwise, if x is an SNaN, return x represented as a QNaN.
Otherwise, if x is a negative, nonzero value, return the default QNaN.
Otherwise, return an estimate of the reciprocal of the square root of x having unbounded exponent range.
reset_xflags()
vxsnan_flag is set to 0.
vximz_flag is set to 0.
vxidi_flag is set to 0.
vxisi_flag is set to 0.
vxzdz_flag is set to 0.
vxsqrt_flag is set to 0.
vxcvi_flag is set to 0.
vxvc_flag is set to 0.
ox_flag is set to 0.
ux_flag is set to 0.
xx_flag is set to 0.
zx_flag is set to 0.

```

\section*{Version 3.0}

\section*{RoundToDP( \(\mathrm{x}, \mathrm{y}\) )}
\(x\) is a 2-bit unsigned integer specifying one of four rounding modes.
\begin{tabular}{ll} 
Ob00 & Round to Nearest Even \\
Ob01 & Round towards Zero \\
Ob10 & Round towards + Infinity \\
Ob11 & Round towards - Infinity
\end{tabular}
y is a normalized floating-point value having unbounded range and precision.
Return the value \(y\) rounded to double-precision under control of the rounding mode specified by \(x\).
```

if I sQNaN(y) then return ConvertFPtoDP(y)
if IsInf(y) then return ConvertFPtoDP(y)
if IsZero(y) then return ConvertFPtoDP(y)
if y<Nmin then do
if UE=0 then do
if x=0b00 then r \& RoundToDPNearEven( DenormDP(y) )
if x=ObOl then r \& RoundToDPTrunc( DenormDP(y) )
if x=0b10 then r \& RoundToDPCeil( DenormDP(y) )
if x=0b11 then r \& RoundToDPFIoor( DenormDP(y) )
ux_flag}\leftarrowx\mp@subsup{x}{_}{\prime}fla
return(ConvertFPtoDP(r))
end
else do
y}\leftarrow\mathrm{ Scalb (y,+1536)
ux_flag}\leftarrow
end
end
if x=0bOO then r \& RoundToDPNearEven(y)
if x=0bO1 then r \& RoundToDPTrunc(y)
if x=0b10 then r \& RoundToDPCeil(y)
if x=0b11 then r \& RoundToDPFloor(y))
if r>Nmax then do
if OE=0 then do
if x=0b00 then r \& sign ? -Inf: tInf
if x=0b01 then r \& sign ? - Nmax : +Nmax
if x=0blo then r \& sign ? -Nmax : +Inf
if x=Obll then r \& sign ? -Inf : +Nmax
ox flag}\leftarrow0b
xx_flag}\leftarrow0b
inc_flag}\leftarrow0b
return(ConvertFPtoDP(r))
end
else do
r}\leftarrow\mathrm{ Scalb(r,-1536)
0x_flag}\leftarrow
end
end
return(ConvertFPtoDP(r))

```

\section*{RoundToDPCeil(x)}
\(x\) is a floating-point value having unbounded range and precision.
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if x is an Infinity, return x .
Otherwise, do the following.
Return the smallest floating-point number having unbounded exponent range but double-precision significand precision that is greater or equal in value to \(x\).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x \_f l a g\) is set to 1 .

\section*{RoundToDPFloor(x)}
\(x\) is a floating-point value having unbounded range and precision.
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if x is an Infinity, return x .
Otherwise, do the following.
Return the largest floating-point number having unbounded exponent range but double-precision significand precision that is lesser or equal in value to \(x\).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x \_f l a g\) is set to 1 .

\section*{RoundToDPIntegerCeil(x)}
\(x\) is a double-precision floating-point value.
If \(x\) is an \(S N a N, v x\) snan_flag is set to 1 .
If x is a QNaN , return x .
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if x is an infinity, return x .
Otherwise, do the following.
Return the smallest double-precision floating-point integer value that is greater or equal in value to \(x\).
If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x \_f l a g\) is set to 1 .

\section*{RoundToDPIntegerFloor(x)}
x is a double-precision floating-point value.
If \(x\) is an \(S N a N\), vxsnan_flag is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if x is an SNaN , return x represented as a QNaN .
Otherwise, if x is an infinity, return x .
Otherwise, do the following.
Return the largest double-precision floating-point integer value that is lesser or equal in value to \(x\)
If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x_{-} f l a g\) is set to 1 .

\section*{RoundToDPIntegerNearAway(x)}
\(x\) is a double-precision floating-point value.
If \(x\) is an \(\mathrm{SNaN}, \mathrm{vxsnan} \_f l a g\) is set to 1 .
If x is a QNaN , return x .
Otherwise, if x is an SNaN , return x represented as a QNaN .
Otherwise, if x is an infinity, return x .
Otherwise, do the following.
Return the largest double-precision floating-point integer value that is lesser or equal in value to \(x+0.5\) if \(x>0\), or the smallest double-precision floating-point integer that is greater or equal in value to \(x-0.5\) if \(x<0\).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x \_f l a g\) is set to 1 .

\section*{RoundToDPIntegerNearEven(x)}
\(x\) is a double-precision floating-point value.
If \(x\) is an \(\mathrm{SNaN}, \mathrm{vxsnan} \_f l a g\) is set to 1 .
If x is a QNaN , return x .
Otherwise, if x is an SNaN , return x represented as a QNaN .
Otherwise, if x is an infinity, return x .
Otherwise, do the following.
Return the double-precision floating-point integer value that is nearest in value to x (in case of a tie, the double-precision floating-point integer value with the least-significant bit equal to 0 is used).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x \_f l a g\) is set to 1 .

\section*{RoundToDPIntegerTrunc(x)}
\(x\) is a double-precision floating-point value.
If \(x\) is an \(S N a N, v \times s n a n+f l a g\) is set to 1 .
If x is a QNaN , return x .
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if x is an infinity, return x .
Otherwise, do the following.
Return the largest double-precision floating-point integer value that is lesser or equal in value to \(x\) if \(x>0\), or the smallest double-precision floating-point integer value that is greater or equal in value to \(x\) if \(x<0\).

If the magnitude of the value returned is greater than \(x, i n c_{-} f l a g\) is set to 1 .
If the value returned is not equal to \(x, x_{x_{-}} f \mid a g\) is set to 1 .
RoundToDPNearEven(x)
\(x\) is a floating-point value having unbounded range and precision.
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an Infinity, return \(x\).
Otherwise, do the following.
Return the floating-point number having unbounded exponent range but double-precision significand precision that is nearest in value to \(x\) (in case of a tie, the floating-point number having unbounded exponent range but double-precision significand precision with the least-significant bit equal to 0 is used).

If the magnitude of the value returned is greater than \(x, i n c \_f l a g\) is set to 1 .
If the value returned is not equal to \(x, x_{x_{-}} f \mid a g\) is set to 1 .
RoundToDPTrunc(x)
\(x\) is a floating-point value having unbounded range and precision.
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if x is an Infinity, return x .
Otherwise, do the following.
Return the largest floating-point number having unbounded exponent range but double-precision significand precision that is lesser or equal in value to \(x\) if \(x>0\), or the smallest floating-point number having unbounded exponent range but double-precision significand precision that is greater or equal in value to \(x\) if \(x<0\).

If the magnitude of the value returned is greater than \(x, i n c \_f a g\) is set to 1 .
If the value returned is not equal to \(x, x_{x_{-}} f \mid a g\) is set to 1 .

\section*{Version 3.0}

\section*{RoundToSP( \(x, y\) )}
\(x\) is a 2-bit unsigned integer specifying one of four rounding modes.
\begin{tabular}{ll} 
Ob00 & Round to Nearest Even \\
Ob01 & Round towards Zero \\
Ob10 & Round towards + Infinity \\
Ob11 & Round towards - Infinity
\end{tabular}
y is a normalized floating-point value having unbounded range and precision.
Return the value \(y\) rounded to single-precision under control of the rounding mode specified by \(x\).
```

f I SQNaN(y) then return ConvertFPtoSP(y)
if IsInf(y) then return ConvertFPtoSP(y)
if IsZero(y) then return ConvertFPtoSP(y)
if y<Nmin then do
if UE=0 then do
if x=0bOO then r \& RoundToSPNearEven( DenormSP(y) )
if x=ObOl then r \& RoundToSPTrunc( DenormSP(y) )
if x=0b10 then r \& RoundToSPCeil( DenormSP(y) )
if x=0b11 then r \& RoundToSPFIoor( DenormSP(y) )
ux_flag}\leftarrow <x_flag
return(ConvertFPtoSP(r))
end
else do
y}\leftarrow\mathrm{ Scalb(y,+192)
ux_flag}\leftarrow
end
end
if x=0bOO then r \& RoundToSPNearEven(y)
if x=0bO1 then r \& RoundToSPTrunc(y)
if x=0b10 then r \& RoundToSPCeil(y)
if x=0b11 then r \& RoundToSPFIoor(y))
if r>Nmax then do
if OE=0 then do
if x=0b00 then r \& sign ? -Inf: tInf
if x=0b01 then r \& sign ? - Nmax : +Nmax
if x=0blo then r \& sign ? -Nmax : +Inf
if x=Obll then r s sign ? -Inf: +Nmax
ox_flag}\leftarrow0b
xx_flag}\leftarrow0b
inc_flag}\leftarrow0b
return(ConvertFPtoSP(r))
end
else do
r}\leftarrow\mathrm{ Scalb(r,-192)
0x_flag}\leftarrow
end
end
return(ConvertFPtoSP(r))

```

\section*{RoundToSPCeil(x)}
\(x\) is a floating-point value having unbounded range and precision.
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if x is an Infinity, return x .
Otherwise, do the following.
Return the smallest floating-point number having unbounded exponent range but single-precision significand precision that is greater or equal in value to \(x\).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x \_f l a g\) is set to 1 .
RoundToSPFloor(x)
\(x\) is a floating-point value having unbounded range and precision.
If x is a QNaN , return x .
Otherwise, if x is an Infinity, return x .
Otherwise, do the following.
Return the largest floating-point number having unbounded exponent range but single-precision significand precision that is lesser or equal in value to \(x\).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x \_f l a g\) is set to 1 .

\section*{RoundToSPIntegerCeil(x)}
\(x\) is a single-precision floating-point value.
If \(x\) is an \(S N a N, v x s n a n \_f l a g\) is set to 1 .
If x is a QNaN , return x .
Otherwise, if x is an SNaN , return x represented as a QNaN .
Otherwise, if x is an infinity, return x .
Otherwise, do the following.
Return the smallest single-precision floating-point integer value that is greater or equal in value to \(x\).
If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x \_f l a g\) is set to 1 .

\section*{Version 3.0}

\section*{RoundToSPIntegerFloor(x)}
\(x\) is a single-precision floating-point value.
If \(x\) is an \(S N a N, v x\) snan_flag is set to 1 .
If x is a QNaN , return x .
Otherwise, if x is an SNaN , return x represented as a QNaN .
Otherwise, if x is an infinity, return x .
Otherwise, do the following.
Return the largest single-precision floating-point integer value that is lesser or equal in value to x .
If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x_{-} f l a g\) is set to 1 .

\section*{RoundToSPIntegerNearAway(x)}
\(x\) is a single-precision floating-point value.
If \(x\) is an \(\mathrm{SNaN}, \mathrm{vxsnan} \_f l a g\) is set to 1 .
If x is a QNaN , return x .
Otherwise, if x is an SNaN , return x represented as a QNaN .
Otherwise, if x is an infinity, return x .
Otherwise, do the following.
Return x if x is a floating-point integer; otherwise return the largest single-precision floating-point integer value that is lesser or equal in value to \(x+0.5\) if \(x>0\), or the smallest single-precision floating-point integer value that is greater or equal in value to \(x-0.5\) if \(x<0\).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x_{\_} f l a g\) is set to 1 .

\section*{RoundToSPIntegerNearEven(x)}
\(x\) is a single-precision floating-point value.
If \(x\) is an \(\mathrm{SNaN}, \mathrm{vxsnan} \_f l a g\) is set to 1 .
If x is a QNaN , return x .
Otherwise, if x is an SNaN , return x represented as a QNaN .
Otherwise, if \(x\) is an infinity, return \(x\).
Otherwise, do the following.
Return x if x is a floating-point integer; otherwise return the single-precision floating-point integer value that is nearest in value to \(x\) (in case of a tie, the single-precision floating-point integer value with the least-significant bit equal to 0 is used).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x \_f l a g\) is set to 1 .

\section*{RoundToSPIntegerTrunc(x)}
\(x\) is a single-precision floating-point value.
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if x is an SNaN , return x represented as a QNaN, and vxsnan_flag is set to 1 .
Otherwise, if x is an infinity, return x .
Otherwise, do the following.
Return the largest single-precision floating-point integer value that is lesser or equal in value to \(x\) if \(x>0\), or the smallest single-precision floating-point integer value that is greater or equal in value to x if \(\mathrm{x}<0\).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x_{-} f l a g\) is set to 1 .

\section*{RoundToSPNearEven(x)}
\(x\) is a floating-point value having unbounded range and precision.
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if x is an Infinity, return x .
Otherwise, do the following.
Return the floating-point number having unbounded exponent range but single-precision significand precision that is nearest in value to \(x\) (in case of a tie, the floating-point number having unbounded exponent range but single-precision significand precision with the least-significant bit equal to 0 is used).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x_{x} f l a g\) is set to 1 .

\section*{RoundToSPTrunc( \(x\) )}
\(x\) is a floating-point value having unbounded range and precision.
If x is a QNaN , return x .
Otherwise, if x is an Infinity, return x .
Otherwise, do the following.
Return the largest floating-point number having unbounded exponent range but single-precision significand precision that is lesser or equal in value to \(x\) if \(x>0\), or the smallest single-precision floating-point number that is greater or equal in value to x if \(\mathrm{x}<0\).

If the magnitude of the value returned is greater than \(x\), inc_flag is set to 1 .
If the value returned is not equal to \(x, x x \_f l a g\) is set to 1 .

\section*{Scalb( \(x, y\) )}
\(x\) is a floating-point value having unbounded range and precision.
y is a signed integer.
Result of multiplying the floating-point value \(x\) by \(2^{y}\).

\section*{SetFX(x)}
\(x\) is one of the exception flags in the FPSCR.
If the contents of \(x\) is \(0, F X\) and \(x\) are set to 1 .

\section*{Version 3.0}

SquareRootDP(x)
x is a double-precision floating-point value.
If \(x\) is an \(S N a N, v x s n a n \_f l a g\) is set to 1 .
If \(x\) is a negative, nonzero value, vxsqrt_flag is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN .
Otherwise, if \(x\) is a negative, nonzero value, return the default QNaN .
Otherwise, return the normalized square root of \(x\), having unbounded range and precision.
SquareRootSP(x)
x is a single-precision floating-point value.
If \(x\) is an \(\mathrm{SNaN}, \mathrm{vxsnan} \_f l a g\) is set to 1 .
If \(x\) is a negative, nonzero value, vxsqrt_flag is set to 1 .
If \(x\) is a \(Q N a N\), return \(x\).
Otherwise, if \(x\) is an SNaN , return x represented as a QNaN.
Otherwise, if \(x\) is a negative, nonzero value, return the default QNaN .
Otherwise, return the normalized square root of \(x\), having unbounded range and precision.

\subsection*{7.6.3 VSX Instruction Descriptions}
\begin{tabular}{l} 
Load VSX Scalar Doubleword DS-form \\
Ixsd \\
\begin{tabular}{|l|l|l|l|ll|l|}
\hline 0 & 57 & 6 & VRT & RA
\end{tabular} \\
\hline 0
\end{tabular} \begin{tabular}{l}
11
\end{tabular}

\footnotetext{
if MSR. VEC=0 then Vector_Unavailablel)
\(E A \leftarrow((R A=0) ? 0: G P R[R A])+E X T S(D S) \ll 2\)
VSR[VRT +32\(]\). dwor d \([0] \leftarrow M E M(E A, 8)\)
VSR[VRT+32].dword[1] \& OxUUUU_UUUU_UUUU_UUUU
}

Let \(X T\) be the value VRT +32 .
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if \(R A=0\), and the signed integer value \(D S \ll 2\).

When Big-Endian byte ordering is employed, the contents of the doubleword in storage at address EA are placed into load_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte 0 ofload_data,
- the contents of the byte in storage at address EA +1 are placed into byte 1 of load_data, and so forth until
- the contents of the byte in storage at address EA +7 are placed into byte 7 of load_data.

When Little-Endian byte ordering is employed, let load_dat a be the contents of the doubleword in storage at address EA such that;
- the contents of the byte in storage at address EA are placed into byte 7 of load_data,
- the contents of the byte in storage at address EA+1 are placed into byte 6 of load_data, and so forth until
- the contents of the byte in storage at address EA +7 are placed into byte 0 of load_data.
load_data is placed into doubleword element 0 of VSR[ XT].

The contents of doubleword element 1 of VSR[XT] are undefined.

\section*{Special Registers Altered: \\ None}

\section*{Load VSX Scalar Doubleword Indexed X-form} lxsdx XT,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline - 31 & 6 & T & 11 & RA & 16 & RB & 21 & 588 & \(\left\lvert\, \begin{aligned} & T X \\ & 31\end{aligned}\right.\) \\
\hline
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable()
EA\leftarrow((RA=0)? 0:GP[[RA]) +GPR[RB]
VSR[32xTX+T], dword[0] \& MEM[EA,8)
VSR[32xTX+T], dwor d[1] \& OxUUUU_UUUU_UUUU_UUUU

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

When Big-Endian byte ordering is employed, the contents of the doubleword in storage at address EA are placed into load_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte 0 of load_data,
- the contents of the byte in storage at address \(E A+1\) are placed into byte 1 of load_data, and so forth until
- the contents of the byte in storage at address EA +7 are placed into byte 7 of load_data.

When Little-Endian byte ordering is employed, the contents of the doubleword in storage at address EA are placed into load_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte 7 of 10 ad_data,
- the contents of the byte in storage at address EA +1 are placed into byte 6 of load_data, and so forth until
- the contents of the byte in storage at address EA +7 are placed into byte 0 of load_data.
load_data is placed into doubleword element 0 of VSR[XT].

The contents of doubleword element 1 of VSR[ XT] are undefined.

\section*{Special Registers Altered \\ None}

\section*{Version 3.0}
\begin{tabular}{l}
\hline VSR Data Layout for Ixsd \\
tgt = VSR[ XT] \\
\begin{tabular}{|c|lr|}
\hline MEM(EA,8) & undefined \\
\hline 0 & 64 & \\
\hline VSR Data Layout for Ixsdx \\
t gt = VSR[ XT] \\
\begin{tabular}{|c|c|}
\hline \multicolumn{3}{|c|}{ MEM(EA,8) } & undefined \\
\hline 0 & 64
\end{tabular} \\
\hline
\end{tabular}
\end{tabular}

Load VSX Scalar as Integer Byte \& Zero Indexed X-form
Ixsibzx
\begin{tabular}{|c|cc|c|c|c|c|c|}
\hline 31 & 6 & T & RA, RA, RB \\
0 & 31 & 6 & & RA & RB & & 781 \\
31
\end{tabular}
if \(T X=0\) \& MSR.VSX=O then VSX_Unavailablell
if \(T X=1\) \& MSR.VEC=0 then Vector_Unavailable()
\(E A \leftarrow((R A=0) ? 0: G P R[R A])+G P R[R B]\)
\(\operatorname{VSR}[32 \times T X+T]\).dword \([0] \leftarrow\) EXTZ64 \((\) MEM \((E A, 1))\)
VSR[32xTX + T] . dword \([1] \leftarrow\) OxUUUU UUUU UUUU UUUU
Let XT be the value \(32 \times T X+T\).
Let the effective address (EA) be sum of the contents of \(G P R[R A]\), or 0 if RA is equal to 0 , and the contents of GPR[RB].

The unsigned integer in the byte in storage addressed by EA is placed in doubleword element 0 of VSR[XT]. The contents of doubleword element 1 of VSR[ XT] are undefined.

\section*{Special Registers Altered:}

None

\section*{Load VSX Scalar as Integer Halfword \& Zero Indexed X-form}
Ixsihzx
XT,RA,RB
\begin{tabular}{|c|cc|c|c|cc|c|}
\hline 31 & 6 & T & & RA & & RB & \\
\hline 0 & 813 & \\
31
\end{tabular}
if TX=O \& MSR. VSX=O then VSX_Unavailablel)
if \(T X=1\) \& MSR.VEC=O then Vector_Unavailablel)
\(E A \leftarrow((R A=0) ? 0: G P R[R A])+G P R[R B]\)
VSR[32xTXXT]. dwor d[0] \(\leftarrow \operatorname{EXTZ64(MEM(EA,2))}\)
VSR[32xTXTT], dwor d [1] \& OxUUUU_ UUUU_ UUUU_UUUU
Let XT be the value \(32 \times T X+T\).
Let the effective address (EA) be sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

The unsigned integer in the halfword in storage addressed by EA is placed in doubleword element 0 of VSR[XT]. The contents of doubleword element 1 of VSR[XT] are undefined.

Special Registers Altered:
None
VSR Data Layout for Ixsihzx
tgt = VSR[XT]


Load VSX Scalar as Integer Word Algebraic Indexed \(X\)-form
Ixsiwax
\begin{tabular}{|c|cc|c|c|c|c|}
\hline 0 & 31 & XT,RA,RB \\
0 & T & 6 & RA & \({ }_{11}\) & RB & 21 \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablell
\(E A \leftarrow((R A=0) ? 0: G P R[R A])+G P R[R B]\)
\(\operatorname{VSR}[32 \times T X+T]\). dword \([0] \leftarrow \operatorname{EXTS} 64(\) MEM \((E A, 4))\)
\(\operatorname{VSR}[32 \times T X+T]\). dwor d \([1] \leftarrow\) OXUUUU_UUUU_UUUU_UUUU
Let \(X T\) be the value \(32 \times T X+T\).
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

When Big-Endian byte ordering is employed, the contents of the word in storage at address EA are placed into load_dat a in such an order that;
- the contents of the byte in storage at address EA are placed into byte 0 of load_data,
- the contents of the byte in storage at address EA+1 are placed into byte 1 of load_data,
- the contents of the byte in storage at address EA +2 are placed into byte 2 of 10 od_data, and
- the contents of the byte in storage at address EA +3 are placed into byte 3 of load_data.

When Little-Endian byte ordering is employed, the contents of the word in storage at address EA are placed into load_dat a in such an order that;
- the contents of the byte in storage at address EA are placed into byte 3 of load_data,
- the contents of the byte in storage at address EA +1 are placed into byte 2 of load_data,
- the contents of the byte in storage at address EA +2 are placed into byte 1 of \(\mid\) oad_dat a, and
- the contents of the byte in storage at address EA +3 are placed into byte 0 of load_data.
load_data is sign-extended to a doubleword and placed in doubleword element 0 of VSR[ XT].

The contents of doubleword element 1 of VSR[ XT] are undefined.

\section*{Special Registers Altered}

None

\section*{VSR Data Layout for Ixsiwax}
tgt \(=\) VSR[ XT]
\begin{tabular}{|l|l|}
\hline & \multicolumn{1}{c|}{ undefined } \\
\hline 0 & 64
\end{tabular}

Load VSX Scalar as Integer Word and Zero Indexed X-form
Ixsiwzx
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & XT,RA,RB \\
0 & & 6 & & T & RA & \\
\hline 11 & RB & & 12 & \begin{tabular}{|c|} 
TX \\
31
\end{tabular} \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablell
\(E A \leftarrow((R A=0) ? 0: G P R[R A])+G P R[R B]\)
\(\operatorname{VSR}[32 \times T X+T]\), dword \([0] \leftarrow\) ExtendZero(MEM(EA, 4))
VSR[32xTX+T]. dwor \(d[1] \leftarrow\) OxUUUU_UUUU_UUUU_UUUU
Let \(X T\) be the value \(32 \times T X+T\).
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

When Big-Endian byte ordering is employed, the contents of the word in storage at address EA are placed intoload_dat a in such an order that;
- the contents of the byte in storage at address EA are placed into byte 0 of load_data,
- the contents of the byte in storage at address EA +1 are placed into byte 1 of load_data,
- the contents of the byte in storage at address EA +2 are placed into byte 2 of load_data, and
- the contents of the byte in storage at address EA +3 are placed into byte 3 of load_data.

When Little-Endian byte ordering is employed, the contents of the word in storage at address EA are placed into load_dat a in such an order that;
- the contents of the byte in storage at address EA are placed into byte 3 of load_data,
- the contents of the byte in storage at address EA +1 are placed into byte 2 of load_data,
- the contents of the byte in storage at address EA +2 are placed into byte 1 of load_data, and
- the contents of the byte in storage at address EA +3 are placed into byte 0 of load_data.
load_data is zero-extended to a doubleword and placed in doubleword element 0 of VSR[ XT] .

The contents of doubleword element 1 of VSR[ XT] are undefined.

\section*{Special Registers Altered}

None

VSR Data Layout for Ixsiwzx
tgt = VSR[XT]
\begin{tabular}{|l|l|l|}
\hline \(0 \times 0000 \_0000\) & word[1] & undefined \\
\hline 0 & 64 & 127 \\
\hline
\end{tabular}

Load VSX Scalar Single DS-form
Ixssp
\begin{tabular}{|c|c|c|cc|c|}
\hline 57 & VRT & RA & & DS & 3 \\
0 & 6 & 11 & 16 & & 3031 \\
\hline
\end{tabular}
if MSR. VEC=0 then Vector_Unavailablell
\(E A \leftarrow((\) RA \(=0) ? 0:\) GPR[RA] \()+\) EXTS (DS \(\mid\) ODOO \()\)

VSR[VRT +32]. dwor d[ 1] \& OxUUUU_UUUZ_UUUU_UUUU
Let \(X T\) be the value VRT +32 .
Let EA be the sum of the contents of GPR[RA], or 0 if \(R A=0\), and the signed integer value \(D S \| O b 00\).

When Big-Endian byte ordering is employed, the contents of the word in storage at address EA are placed into load_dat a in such an order that;
- the contents of the byte in storage at address EA are placed into byte 0 of load_data,
- the contents of the byte in storage at address EA+1 are placed into byte 1 of load_data,
- the contents of the byte in storage at address EA+2 are placed into byte 2 of load_data, and
- the contents of the byte in storage at address EA+3 are placed into byte 3 of load_data.

When Little-Endian byte ordering is employed, the contents of the word in storage at address EA are placed into load_dat a in such an order that;
- the contents of the byte in storage at address EA are placed into byte 3 of load_data,
- the contents of the byte in storage at address EA +1 are placed into byte 2 of load_data,
- the contents of the byte in storage at address EA+2 are placed into byte 1 of 1 oad_data, and
- the contents of the byte in storage at address EA+3 are placed into byte 0 of load_data.
load_data, interpreted as a single-precision floating-point value, is placed into doubleword element 0 of VSR[ VRT +32] in double-precision format.

The contents of doubleword element 1 of VSR[ VRT +32] are undefined.

\section*{Special Registers Altered: None}

\section*{Load VSX Scalar Single-Precision Indexed X-form}
Ixsspx
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & & T & RA & RB \\
\hline 0 & & 6 & & 11 & 16 & \\
\hline 1 & 524 & TX \\
31
\end{tabular}
if MSR. VSX=0 then VSX_Unavailablel)
\(E A \leftarrow((R A=0) ? 0: G P R[R A])+G P R[R B]\)

VSR[ VRT + 32 ]. dwor d [ 1 ] \& OxUUUU_UUUU_UUUU_ UUUU
Let \(X T\) be the value \(32 \times T X+T\).
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of \(G P R[R B]\).

When Big-Endian byte ordering is employed, the contents of the word in storage at address EA are placed intoload_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte 0 of 1 oad_data,
- the contents of the byte in storage at address \(E A+1\) are placed into byte 1 of 10 ad_data,
- the contents of the byte in storage at address EA +2 are placed into byte 2 of Ioad_data, and
- the contents of the byte in storage at address EA +3 are placed into byte 3 of load_data.

When Little-Endian byte ordering is employed, the contents of the word in storage at address EA are placed intoload_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte 3 of 1 oad_data,
- the contents of the byte in storage at address \(E A+1\) are placed into byte 2 of 10 ad_data,
- the contents of the byte in storage at address EA +2 are placed into byte 1 of 1 oad_data, and
- the contents of the byte in storage at address EA +3 are placed into byte 0 of load_data.
load_data, interpreted as a single-precision floating-point value, is placed in doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[ XT] are undefined.

\section*{Special Registers Altered None}
```

VSR Data Layout for Ixssp
tgt = VSR[XT]

|  | undefined |
| :--- | :--- |
| 0 | 64 |

VSR Data Layout for Ixsspx
tgt = VSR[ XT]
|DP

```

\section*{Load VSX Vector Byte*16 Indexed X-form}
Ixvb16x
\begin{tabular}{|c|cc|c|c|c|c|}
\hline 31, RA, RB \\
0 & 31 & 6 & T & 11 & RA & \({ }_{16}\) \\
\hline
\end{tabular}
if \(T X=0\) \& MSR.VSX=O then VSX_Unavailable(l)
if \(T X=1\) \& MSR.VEC=O then Vector Unavailablell
\(E A \leftarrow R A=0 ? G P R[R B]: G P R[R A]+G P R[R B]\)
do \(\mathrm{i}=0\) to 15
\(\operatorname{VSR}[32 \times T X+T]\), byte \([i] \leftarrow \operatorname{MEM}(E A+i, 1)\)
end
Let \(X T\) be the value \(32 \times T X+T\).
Let the effective address (EA) be the sum of the contents of \(G P R[R A]\), or 0 if \(R A\) is equal to 0 , and the contents of GPR[RB].

For each integer value from 0 to 15 , do the following.
The contents of the byte in storage at address \(E A+i\) are placed into byte element \(i\) of VSR[ XT] ,

\section*{Special Registers Altered:}

None

\section*{Example: Loading data using Load VSX Vector} Byte*16 Indexed


Big-endian storage image of \(X\)


Loading a vector of 16 byte elements from Big-Endian storage in VSR[ XT] using Ixvb16x, retaining left-to-right element ordering.
```


# Assumptions

# GPR[PX] = address of X

Ixvb16x xX,ro,rPX

```
    VSR1 W] : \begin{tabular}{|l|l|l|l|l|l|l|l|l|l|l|l|l|l|l|l|}
\hline F0 & F1 & F2 & F3 & F4 & F5 & F6 & F7 & E0 & E1 & E2 & E3 & E4 & E5 & E6 & E7 \\
\hline
\end{tabular}

Loading a vector of 16 byte elements from Little-Endian storage in VSR[XT] using Ixvb16x, retaining left-to-right element ordering.
```


# Assumptions

# GPR[PX] = address of X

Ixvb16x xX,r0,rPX
VSR[X]:

```


\section*{Load VSX Vector Doubleword*2 Indexed} X-form
Ixvd2x \begin{tabular}{l} 
XT,RA,RB \\
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & 31 & T & RA & & RB & \\
0 & & 6 & & 11 & & 16
\end{tabular} \\
\hline
\end{tabular}
if MSR.VSX=O then VSX_Unavailable()
\(E A \leftarrow R A=0 ? G P R[R B]: G P R[R A]+G P R[R B]\)
\(\operatorname{VSR}[32 x T X+T]\). dword \([0] \leftarrow\) MEM \((E A, 8)\)
\(\operatorname{VSR}[32 \times T X+T]\). dword \([1] \leftarrow \operatorname{MEM}(E A+8,8)\)
Let \(X T\) be the value \(32 \times T X+T\).
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[ RB].

For each integer value i from 0 to 1 , do the following. When Big-Endian byte ordering is employed, the contents of the doubleword in storage at address EA \(+8 \times\) i are placed intoload_data in such an order that;
- the contents of the byte in storage at address \(E A+8 \times i\) are placed into byte element 0 of load_data,
- the contents of the byte in storage at address \(E A+8 \times i+1\) are placed into byte element 1 of load_dat a, and so forth until
- the contents of the byte in storage at address \(E A+8 \times i+7\) are placed into byte element 7 of load_data.

When Little-Endian byte ordering is employed, the contents of the doubleword in storage at address EA \(+8 \times\) i are placed intoload_data in such an order that;
- the contents of the byte in storage at address \(E A+8 x i\) are placed into byte element 7 of load_data,
- the contents of the byte in storage at address \(E A+8 \times i+1\) are placed into byte element 6 of load_dat a, and so forth until
- the contents of the byte in storage at address \(E A+8 \times i+7\) are placed into byte element 0 of load_data.
load_data is placed into doubleword element i of VSR[XT].

VSR Data Layout for Ixvd2x
tgt \(=\) VSR[ XT]
\begin{tabular}{|l|l|}
\hline &.\(d\) word[0] \\
\hline 0 & \(. \quad . d\) word [1] \\
\hline
\end{tabular}
\begin{tabular}{llll}
\hline \multicolumn{2}{l}{ Extended Mnemonic } & \multicolumn{2}{l}{ Equivalent To } \\
\hline Ixvx & \(\mathrm{XT}, \mathrm{RA}, \mathrm{RB}\) & Ixvd 2 x & \(\mathrm{XT}, \mathrm{RA}, \mathrm{RB}\)
\end{tabular}

Usage: The Ixvx extended mnemonic should be be I used for vector load operations when using Big-Endian byte-ordering, independent of element size.

\section*{Special Registers Altered}

None

Load VSX Vector with Length X-form
\begin{tabular}{|c|c|c|c|c|}
\hline \multicolumn{5}{|l|}{Ixvl XT,RA,RB} \\
\hline \[
31
\] & \[
{ }_{6} \mathrm{~T}
\] & \[
{ }_{11} \mathrm{RA}
\] & \({ }_{16} \mathrm{RB}\) & \[
26
\] \\
\hline \multicolumn{5}{|l|}{\multirow[t]{2}{*}{\begin{tabular}{l}
if \(T X=0\) \& MSR. VSX \(=0\) then VSX_Unavailable(l) \\
if \(T X=1\) \& MSR.VEC=0 then Vector Unavailable(l)
\end{tabular}}} \\
\hline & & & & \\
\hline \multicolumn{5}{|l|}{\multirow[t]{2}{*}{\[
\begin{aligned}
& E A \leftarrow(R A=0) ? 0: \text { GPR[RA] } \\
& n b \leftarrow \text { EXTZ(GPR[RB]. bit }[0: 7])
\end{aligned}
\]}} \\
\hline & & & & \\
\hline \multicolumn{5}{|l|}{load_data \(\leftarrow 0 \times 0000 \_00000^{\prime} 0000 \_0000 \_0000 \_0000 \_0000 \_0000\)} \\
\hline \multicolumn{5}{|l|}{if MSR.LE = 0 then \(/ / \mathrm{Big}\)-Endian byte-ordering load_data,byte[0:nb-1] \(\leftarrow M E M(E A, n b)\)} \\
\hline \multicolumn{5}{|l|}{```
else |l Little-Endian byte-ordering
    load_data,byte[16.nb:15] \leftarrowMEM(EA, nb)
```} \\
\hline \multicolumn{5}{|l|}{\(\operatorname{VSR}[32 \times T X+T] \leftarrow 1\) oad_data} \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let the effective address (EA) be the contents of \(G P R[R A]\), or 0 if \(R A\) is equal to 0 .

Let \(n b\) be the unsigned integer value in bits \(0: 7\) of GPR[RB].

If \(n b\) is equal to 0 , the storage access is not performed and the contents of VSR[ XT] are set to 0 .

Otherwise, when Big-Endian byte-ordering is employed, do the following. If nb less than 16 , the contents of the nb bytes in storage starting at address EA are placed into the leftmost nb bytes of VSR[XT], and the contents of the rightmost \(16 \cdot \mathrm{nb}\) bytes of VSR[XT] are set to \(0 \times 00\).

Otherwise, the contents of the quadword in storage at address EA are placed into VSR[ XT] .

Otherwise, when Little-Endian byte ordering is employed, do the following.

If \(n b\) less than 16, the contents of the nb bytes in storage starting at address EA are placed into the rightmost \(n b\) bytes of VSR[XT] in byte-reversed order, and the contents of the leftmost \(16 \cdot \mathrm{nb}\) bytes of VSR[ XT] are set to \(0 \times 00\).

Otherwise, the contents of the quadword in storage at address EA are placed into VSR[XT] in byte-reversed order.

If the contents of bits \(8: 63\) of GPR[RB] are not equal to 0 , the results are boundedly undefined.

\section*{Special Registers Altered: \\ None}

Example: Loading less than 16-byte data into VSR using IxvI
```

char S[14] = "This is a TEST";
short X[6] = { OXEOE1, OXE2E3, OXE4E5, OXE6E7, OXE8E9, OXEAEB };
binary80 Z = OXFOF1F2F3F4F5F6F7F8F9

```

Loading less than 16-byte data from Big-Endian storage in VSR[ XT] using IxvI.

Big-endian storage image of \(\mathrm{S}, \mathrm{X}, \& \mathrm{Z}\)
\(\operatorname{addr}(5)+0 \times 0000\) :
addr (S) \(+0 \times 0010\)
\(\operatorname{addr}(S)+0 \times 0020:\)
\# Assumptions
\# GPR[NS] \(=14\) (length of \(S\) in \# of bytes)
\# GPR[NX] \(=12\) (Iength of \(X\) in \# of bytes)
\# GPR[NZ] \(=10\) (length of \(Z\) in \# of bytes)
\# GPR[PS] = address of \(S\)


Loading less than 16-byte data from Little-Endian storage in VSR[ XT] using IxvI.

Little-endian storage image of \(S, X, \& Z\)

\# Assumptions
\# GPR[NS] \(=14\) (length of \(S\) in \# of bytes)
\# GPR[NX] \(=12\) (Iength of \(X\) in \# of bytes)
\# GPR[NZ] \(=10\) (Iength of \(Z\) in \# of bytes)
\# GPR[PS] = address of \(S\)
\begin{tabular}{lll} 
add & \(r P X, r P S, r N S\) & \# address of \(X\) \\
add & \(r P Z, r P X, r N X\) & \# address of \(Z\) \\
s|di & \(r L S, r N S, 56\) & \\
s|di & \(r L X, r N X, 56\) & \\
\(s \mid d i\) & \(r L Z, r N Z, 56\) & \\
\(|x v|\) & \(x S, r P S, r L S\) & \\
\(|x v|\) & \(X X, r P X, r L X\) & \\
\(|x v|\) & \(x Z, r P Z, r L Z\) &
\end{tabular}

VSR register image of \(S, X, \& Z\)
VSR[S]:
VSR[X]:


\section*{Load VSX Vector Left-justified with Length X-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|l|}{Ixvll XT,RA,RB} \\
\hline \[
\begin{array}{ll} 
& 31 \\
0
\end{array}
\] & \[
\begin{array}{ll} 
& \mathrm{T} \\
6 &
\end{array}
\] & \[
{ }_{11} \mathrm{RA}
\] & \[
{ }_{16} \mathrm{RB}
\] & 21 & 301 & TX \\
\hline
\end{tabular}
if \(T X=0\) \& MSR.VSX=0 then VSX_Unavailablel)
if TX=1 \& MSR.VEC=0 then Vector Unavailable(l)
\(E A \leftarrow(R A=0) ? 0: G P R[R A]\)
\(n b \leftarrow \operatorname{EXTZ}(\) GPR \([R B]\), bit \([0: 7])\)
if nb>0 then do \(i=0\) to \(n b \cdot 1\)
VSR[32xTX CT\(]\), byte[i] \(\leftarrow \operatorname{MEM}(E A+i, 1)\)
end
if nb<l6 then do \(i=n b\) to 15 \(\operatorname{VSR}[32 \times T X+T]\), byte[i] \(\leftarrow 0 \times 00\)
end
Let \(X T\) be the value \(32 \times T X+T\).
Let the effective address (EA) be the contents of \(G P R[R A]\), or 0 if \(R A\) is equal to 0 .

Let \(n b\) be the unsigned integer value in bits \(0: 7\) of GPR[RB].

If \(n b\) is equal to 0 , the storage access is not performed and the contents of VSR[ XT] are set to 0 .

Otherwise, do the following.
If nb less than 16, the contents of the nb bytes in storage starting at address EA are placed into the lefttmost \(n b\) bytes of VSR[ XT], and the contents of the rightmost \(16 \cdot \mathrm{nb}\) bytes of VSR[XT] are set to \(0 \times 00\).

Otherwise, the contents of the quadword in storage at address EA are placed into VSR[ XT].

Data is loaded from storage into VSR[XT] in Big-Endian byte ordering (i.e., the byte in storage at address EA is placed into byte element 0 of VSR[XT], the byte in storage at address EAt1 is placed in byte element 1 of VSR[ XT] , and so forth).

If the contents of bits \(8: 63\) of \(G P R[R B]\) are not equal to 0 , the results are boundedly undefined.

\section*{Special Registers Altered:}

None

\section*{Example: Loading less than 16-byte left-justified data}
```

decimal }\quadX=+1234567890123456789
decimal Y = .123456;
decimal }Z=+1004966723510220

```

Loading less than 16 -byte data from storage in VS R [ XT] , left-justified, using IxvII.
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|r|}{Initial state of VSRs X, Y, \& Z} \\
\hline VSR[ \(X\) ]: & FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF \\
\hline VSR[Y]: & FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF \\
\hline VSR[ Z]: & FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF \\
\hline &  \\
\hline
\end{tabular}

Big-endian \& Little-Endian storage image of \(X, Y, \& Z\)

\# Assumptions
\# GPR[NX] \(=10\) (Iength of \(X\) )
\# GPR[NY] \(=4\) (length of \(Y\) )
\# GPR[NZ] \(=9\) (length of \(Z\) )
\# GPR[PX] = address of \(X\)
\# GPR[PY] = address of \(Y=\) address of \(X+10\)
\# GPR[PZ] = address of \(Z=\) address of \(X+10+4\)
|xv|l \(\quad x X, r P X, r N X\)
|xvII XY,rPY,rNY
|xv|| \(x Z, r P Z, r N Z\)
Final state of VSRs \(X, Y, \& Z\)
VSR[X]: \(\quad 0134677890123456789 \mathrm{C} \mid 000000000000\)
VSR [ Y]: 0123456 D 000000000000000000000000
VSR[Z]:


\section*{Load VSX Vector DQ-form}

Ixv XT,DQ(RA)
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 61 & & T & RA & & DQ & \begin{tabular}{l} 
TX \\
\hline 0
\end{tabular} \\
\hline
\end{tabular}
if TX=0 \& MSR. VSX=O then VSX_Unavailable el)
if \(T X=1 \&\) MSR. VEC=0 then Vector Unavailable ()
\(E A \leftarrow R A=0 ? ~ G P R[R B]: ~ G P R[R A]+E X T S(D Q| | O b 0000)\)
\(\operatorname{VSR}[32 x T X+T] \leftarrow \operatorname{MEM}(E A, 16)\)
Let \(X T\) be the value \(32 \times T X+T\).
Let the effective address ( \(E A\) ) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the signed integer value \(\mathrm{DQ} \| \mathrm{ObOOOO}\).

When Big-Endian byte ordering is employed, the contents of the quadword in storage at address EA are placed into load_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte element 0 of load_data,
- the contents of the byte in storage at address EA+1 are placed into byte element 1 of load_data, and so forth until
- the contents of the byte in storage at address EA+15 are placed into byte element 15 of load_data.

When Little-Endian byte ordering is employed, the contents of the quadword in storage at address EA are placed into load_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte element 15 of load_dat a,
- the contents of the byte in storage at address EA +1 are placed into byte element 14 of load_data, and so forth until
- the contents of the byte in storage at address \(E A+15\) are placed into byte element 0 of 1 oad_data.
load_data is placed into VSR[XT].

\section*{Special Registers Altered}

None

\section*{Load VSX Vector Indexed X-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{Ixvx XT,RA,RB} \\
\hline \[
\begin{array}{ll} 
& 31 \\
0
\end{array}
\] & \({ }_{6} \mathrm{~T}\) & 111 RA & \({ }_{16} \mathrm{RB}\) & & \(\mid /{ }_{25}{ }_{26}\) & 12 & \(\left\lvert\, \begin{aligned} & \text { TX } \\ & 31\end{aligned}\right.\) \\
\hline
\end{tabular}
if TX=O \& MSR.VSX=O then VSX_Unavailablell
if \(T X=1\) \& MSR.VEC=O then Vector Unavailable(l)
\(E A \leftarrow R A=0 ? G P R[R B]: G P R[R A]+G P R[R B]\)
\(\operatorname{VSR}[32 \times T X+T] \leftarrow \operatorname{MEM}(E A, 16)\)
Let \(X T\) be the value \(32 \times T X+T\).
Let the effective address ( \(E A\) ) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[ RB].

When Big-Endian byte ordering is employed, the contents of the quadword in storage at address EA are placed intoload_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte element 0 of 10 odd data,
- the contents of the byte in storage at address EA+1 are placed into byte element 1 of load_data, and so forth until
- the contents of the byte in storage at address EA+15 are placed into byte element 15 of load_data.

When Little-Endian byte ordering is employed, the contents of the quadword in storage at address EA are placed into load_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte element 15 of \(\mid\) oad_dat a,
- the contents of the byte in storage at address \(E A+1\) are placed into byte element 14 of load_data, and so forth until
- the contents of the byte in storage at address EA +15 are placed into byte element 0 of load_data.
load_data is placed into VSR[ XT].

\section*{Special Registers Altered:}

None

\section*{Version 3.0}

\section*{Example: Loading data using Load VSX Vector Indexed}
```

char W[16]={OXFO, OXF1, OXF2, OXF3, OXF4, OXF5, OXF6, OXF7, OXEO, OXE1, OXE2, OXE3, OXE4, OXE5, OXE6, OXE7 };
short X[8] ={OXFOF1, OXF2F3, OXF4F5, OXF6F7, OXEOE1, OXE2E3, OXE4E5,OXE6E7 };
float Y[4] = { OXFOF1_F2F3, OXF4F5_F6F7, OXEOE1_E2E3, OXE4E5_E6E7 };
double Z[2] ={OXFOF1_F2F3_F4F5_F6F7,OXEOE1_E2E3_E4E5_E6ET };

```

Loading 16 bytes of data from Big-Endian storage in VSR[ XT] using Ixvx.

Big-endian storage image of \(\mathrm{W}, \mathrm{X}, \mathrm{Y}, \& \mathrm{Z}\)


\section*{\# Assumptions}
\# GPR[PW] = address of \(W\)
\# GPR[PX] \(=\) address of \(X=G P R[P W]+16\)
\# GPR[PY] \(=\) address of \(Y=G P R[P W]+32\)
\# GPR[PZ] \(=\) address of \(Z=G P R[P W]+48\)
Ixvx \(\quad x W_{1} r 0, r P W\)
\(X X, r 0, r P X\)
\(X Y, r O, r P Y\)
\(X Z, r 0, r P Z\)
Final state of VSRs W, X, Y, \& Z
VSR[W]:

VSR[X]: F0 F1 F2 F3 F4 F5 F6 F7 E0 E1 E2 E3 E4 E5 E6 E7
VSR[Y]:
VSR[Z]:

Loading 16 bytes of data from Little-Endian storage in VSR[ XT] using Ixvx.

Little-endian storage image of \(\mathrm{W}, \mathrm{X}, \mathrm{Y}, \& \mathrm{Z}\)
addr(W+0×0000):
addr(W+0×0010):
addr(W+0×0020):
addr (W+0×0030):

\# Assumptions
\# GPR[PW] = address of \(W\)
\# GPR[PX] \(=\) address of \(X=G P R[P W]+16\)
\# GPR[PY] \(=\) address of \(Y=G P R[P W]+32\)
\# GPR[PZ] \(=\) address of \(Z=G P R[P W]+48\)
\begin{tabular}{ll} 
I \(x v x\) & \(x W, r 0, r P W\) \\
I \(x v x\) & \(x X, r 0, r P X\) \\
I \(x v x\) & \(x Y, r 0, r P Y\) \\
I \(x v x\) & \(x Z, r 0, r P Z\)
\end{tabular}

Final state of VSRs W, X, Y, \& Z
VSR[W]:
VSR[X]:
VSR[Y]:
VSR[Z]:

\section*{Load VSX Vector Doubleword \& Splat Indexed X-form}
Ixvdsx
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline XT,RA,RB \\
\hline 31 & & T & RA & RB & & 332 & TX \\
31
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablell
\(E A \leftarrow R A=0 ? G P R[R B]: G P R[R A]+G P R[R B]\)
load_data \(\leftarrow \operatorname{MEM}(E A, 8)\)
\(\operatorname{VSR}[32 \times T X+T]\).dword \([0] \leftarrow\) load_data
VSR[32 \(2 \times X+T]\). dword[1] \(\leftarrow\) load_data
Let XT be the value \(32 \times T X+T\).
Let EA be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

When Big-Endian byte ordering is employed, the contents of the doubleword in storage at address EA are placed into load_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte element 0 of 1 oad_data,
- the contents of the byte in storage at address EA +1 are placed into byte element 1 of load_data, and so forth until
- the contents of the byte in storage at address EA +7 are placed into byte element 7 of \(\mid\) oad_data.

When Little-Endian byte ordering is employed, the contents of the doubleword in storage at address EA are placed into load_data in such an order that;
- the contents of the byte in storage at address EA are placed into byte element 7 of load_data,
- the contents of the byte in storage at address \(E A+1\) are placed into byte element 6 of load_data, and so forth until
- the contents of the byte in storage at address \(E A+7\) are placed into byte element 0 of \(\mid 0 a d \_d a t a\).
load data is copied into each doubleword element of VSR[ XT].

\section*{Special Registers Altered}

None
VSR Data Layout for Ixvdsx
tgt \(=\) VSR[ \(X T]\)
\begin{tabular}{|l|l|}
\hline &.\(d\) word[0] \\
\hline 0 & 64 \\
\hline
\end{tabular}

\section*{Load VSX Vector Halfword*8 Indexed X-form}
Ixvh8x \begin{tabular}{l} 
XT,RA,RB \\
\begin{tabular}{|c|cc|c|c|c|c|}
\hline 31 & 6 & T & RA & & RB & \\
\hline 0 & & & 11 & 812 & \begin{tabular}{|c|}
\(T x\) \\
31
\end{tabular}
\end{tabular}
\end{tabular}\(.\)\begin{tabular}{l}
21 \\
\hline
\end{tabular}
```

if TX=O \& MSR. VSX=0 then VSX_Unavailable|)
if TX=1\& MSR.VEC=O then Vector_Unavailablel)
EA\leftarrowRA=0 ? GPR[ RB]:GPR[ RA] +GPR[ RB]
do i = 0 to 7
VSR[32xTX+T], hword[i] \&MEM[EA+2xi, 2)

```
end

Let \(X T\) be the value \(32 \times T X+T\).
Let the effective address (EA) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

For each integer value from 0 to 7, do the following. When Big-Endian byte ordering is employed, the contents of the halfword in storage at address \(E A+2 \times i\) are placed intoload data in such an order that;
- the contents of the byte in storage at address \(E A+2 \times i\) are placed into byte element 0 of load_data,
- the contents of the byte in storage at address \(E A+2 \times i+1\) are placed into byte element 1 of load_data.

When Little-Endian byte ordering is employed, the contents of the quadword in storage at address EA are placed into VSR[ XT] in such an order that;
- the contents of the byte in storage at address EA+2xi are placed into byte element 1 of load_data,
- the contents of the byte in storage at address \(E A+2 \times i+1\) are placed into byte element 0 of load_data.
load_data is placed into halfword element i of VSR[ \(\bar{X} T]\).

\section*{Special Registers Altered:}

None

\section*{Example: Loading data using Load VSX Vector Halfword*8 Indexed}
```

short X[] ={0\times0001, 0x1011, 0x2021, 0x3031,
0x4041, 0x5051,0\times6061, 0x7071};

```


Loading a vector of 8 halfword elements from Big-Endian storage in VSR[ XT] using Ixvh8x, retaining left-to-right element ordering.
\[
\begin{aligned}
& \text { \# Assumptions } \\
& \text { \# GPR[PX] }=\text { address of } X \\
& \text { I xuh8x } \quad x X, r 0, r P X \\
& \quad \text { VSR[X]: } \begin{array}{|cc|cc|cc|cc|cc|c|cc|cc|}
\hline 00 & 01 & 10 & 11 & 20 & 21 & 30 & 31 & 40 & 41 & 50 & 51 & 60 & 61 & 70 \\
1 & 71 \\
\hline 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & A & B & C & D & E
\end{array} \\
& \hline
\end{aligned}
\]

Loading a vector of 8 halfword elements from Little-Endian storage in VSR[ XT] using Ixvh8x, retaining left-to-right element ordering.
\[
\begin{aligned}
& \text { \# Assumptions } \\
& \text { \# GPR[PX] = address of } X \\
& \text { I xuh8x } \quad x X, r 0, r P X
\end{aligned}
\]

\section*{Load VSX Vector Word*4 Indexed X-form}
Ixvw4x
\begin{tabular}{|c|cc|c|c|c|cc|c|}
\hline 31 & & T & T RA,RB \\
0 & & 6 & & RA & RB & & 780 & \\
\hline 11 & & 16 & & 21 & & \\
31
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablell
\(E A \leftarrow R A=0 ? G P R[R B]: G P R[R A]+G P R[R B]\)
\(\operatorname{VSR}[32 \times T X+T]\), word \([0] \leftarrow \operatorname{MEM}(E A, 4)\)
\(\operatorname{VSR}[32 x T X+T]\), word \([1] \leftarrow \operatorname{MEM}(E A+4,4)\)
\(\operatorname{VSR}[32 x T X+T]\), word [2] \(\leftarrow \operatorname{MEM}(E A+8,4)\)
\(\operatorname{VSR}[32 x T X+T]\), word \([3] \leftarrow \operatorname{MEM}(E A+12,4)\)
Let XT be the value \(32 \times T X+T\).
Let \(E A\) be the sum of the contents of \(G P R[R A]\), or 0 if \(R A\) is equal to 0 , and the contents of GPR[ RB].

For each integer value i from 0 to 3 , do the following. When Big-Endian byte ordering is employed, the contents of the word in storage at address EA \(+4 \times i\) are placed into load_data in such an order that;
- the contents of the byte in storage at address \(E A+4 \times i\) are placed into byte element 0 of load_data,
- the contents of the byte in storage at address \(E A+4 \times i+1\) are placed into byte element 1 of load_data,
- the contents of the byte in storage at address \(E A+4 \times i+2\) are placed into byte element 2 of load_data, and
- the contents of the byte in storage at address \(E A+4 \times i+3\) are placed into byte element 3 of load_data.

When Little-Endian byte ordering is employed, the contents of the word in storage at address EA \(+4 \times \mathrm{i}\) are placed into word element \(i\) of VSR[ XT] in such an order that;
- the contents of the byte in storage at address EA \(+4 \times \mathrm{xi}\) are placed into byte element 3 of load_data,
- the contents of the byte in storage at address \(E A+4 \times i+1\) are placed into byte element 2 of load_data,
- the contents of the byte in storage at address \(E A+4 \times i+2\) are placed into byte element 1 of load_data, and
- the contents of the byte in storage at address \(E A+4 \times i+3\) are placed into byte element 0 of load_data.
load_data is placed into word element i of VSR[ \(\bar{X} T]\).

\section*{Special Registers Altered None}

VSR Data Layout for Ixvw4x
tgt \(=\) VSR[ \(X T]\)


Load VSX Vector Word \& Splat Indexed X-form
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{11}{|l|}{Ixvwsx \(\quad\) XT,RA,RB} \\
\hline 0 & 31 & 6 & T & 11 & RA & 16 & RB & 21 & 364 & TT \\
\hline
\end{tabular}
if \(T X=0\) \& MSR. VSX=O then VSX_Unavailable(l)
if \(T X=1\) \& MSR. VEC=O then Vector Unavailable(l)
\(E A \leftarrow R A=0 \quad ? G P R[R B]: G P R[R A]+G P R[R B]\)
load_data \(\leftarrow M E M(E A, 4)\)
do \(\mathrm{i}=0\) to 3
VSR[32xTX+T]. word[i] \(\quad\) load_data
end
Let \(X T\) be the value \(32 \times T X+T\).
Let the effective address (EA) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

When Big-Endian byte ordering is employed, the contents of the word in storage at address EA are placed into load_dat a in such an order that;
- the contents of the byte in storage at address EA are placed into byte element 0 of 10 od_dat a,
- the contents of the byte in storage at address EA +1 are placed into byte element 1 of 10 od_dat a,
- the contents of the byte in storage at address EA+2 are placed into byte element 2 of load_data, and
- the contents of the byte in storage at address EA +3 are placed into byte element 3 of load_data.

When Little-Endian byte ordering is employed, the contents of the quadword in storage at address EA are placed intoload_dat a in such an order that;
- the contents of the byte in storage at address EA are placed into byte element 3 of load_dat a,
- the contents of the byte in storage at address EA +1 are placed into byte element 2 of 1 oad_dat a,
- the contents of the byte in storage at address EA +2 are placed into byte element 1 of 10 ad_dat a, and
- the contents of the byte in storage at address \(E A+3\) are placed into byte element 0 of 10 od_data.
load_data is copied into each word element of VSR[XT].

\section*{Special Registers Altered:}

None

\section*{Example: Loading data using Load VSX Vector Word \& Splat Indexed}
int \(\quad X=0 \times F 0 F 1\) F2F3;
Big-endian storage image of \(X\)
\(\operatorname{addr}(\mathrm{X}): \quad\) F0 F1 F2 F3 000000000000000000000000


Little-endian storage image of \(X\)
addr (X): F3 F2 F1 F0000000000000000000000000

Loading scalar word data from Big-Endian storage in VSR[ XT] using Ixvwsx.
```


# Assumptions

# GPR[PX] = address of X

I xvwsx xX,r0,rPX

```

Final state of VSR \(X\)
VSR[ X]: F0 F1 F2 F3 F0 F1 F2 F3 F0 F1 F2 F3 F0 F1 F2 F3 \(\begin{array}{lllllllllllllll} & \\ 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & A & B & C & D & E\end{array}\)

Loading scalar word data from Little-Endian storage in VSR[ XT] using Ixvwsx.
```


# Assumptions

# GPR[PX] = address of X

I xvwsx xX,r0,rPX

```

Final state of VSR X
VSR[X]:


Store VSX Scalar Doubleword DS-form
stxsd VRS,DS(RA)
\begin{tabular}{|c|c|c|cc|c|}
\hline 61 & VRS & RA & & DS & 2 \\
\hline 0 & & 60 & \\
\hline
\end{tabular}
if MSR.VEC=O then Vector_Unavailable()
\(E A \leftarrow R A=0 ? G P R[R B]: G P R[R A]+E X T S(D S| | 0 b 00)\)
MEM(EA, 8) \(\leftarrow\) VSR[VRS +32\(]\). dwor dio]

Let XS be the value VRS + 32 .
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if \(R A=0\), and the signed integer value \(D S \ll 2\).

Letstoredata be the contents of doubleword element 0 of VSR[ XS].

When Big-Endian byte ordering is employed, store_data is placed in the doubleword in storage at address EA in such order that;
- byte 0 of store_data is placed into the byte in storage at address EA,
- byte 1 of store_data is placed into the byte in storage at address \(E A+1\), and so forth until
- byte 7 of store_data is placed into the byte in storage at address \(E A+7\).

When Little-Endian byte ordering is employed, store_data is placed in the doubleword in storage at address EA in such order that;
- the contents of byte 7 of doubleword element 0 of VSR[VRS +32 ] are placed into the byte in storage at address EA,
- the contents of byte 6 of doubleword element 0 of VSR[VRS +32 ] are placed into the byte in storage at address EA+1, and so forth until
- the contents of byte 0 of doubleword element 0 of VSR[VRS +32 ] are placed into the byte in storage at address EA +7 .

Special Registers Altered:
None
VSR Data Layout for stxsd
src = VSR[ XS]
\begin{tabular}{|l|l|}
\hline & unused \\
\hline 0 & 64 \\
\hline
\end{tabular}

Store VSX Scalar Doubleword Indexed X-form
stxsdx XS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & 6 & S & RA & RB & & 716
\end{tabular}\(|\)\begin{tabular}{ll} 
SX \\
31
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable\)
EA\leftarrowRA=0 ? GPR[ RB]:GPR[RA] + GPR[Rb]
MEM[EA,8)}\leftarrow\mathrm{ VSR[XS].dwor d[0]

```

Let \(X S\) be the value \(32 \times 5 X+S\).
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of \(G P R[R B]\).

Let storedata be the contents of doubleword element 0 of VSR[ XS].

When Big-Endian byte ordering is employed, store_data is placed in the doubleword in storage at address EA in such order that;
- byte 0 of store_data is placed into the byte in storage at address EA,
- byte 1 of store_data is placed into the byte in storage at address EA +1 , and so forth until
- byte 7 of store_data is placed into the byte in storage at address \(\mathrm{EA}+7\).

When Little-Endian byte ordering is employed, store_data is placed in the doubleword in storage at address EA in such order that;
- byte 0 of store_data is placed into the byte in storage at address \(\mathrm{EA}+7\),
- byte 1 of store_data is placed into the byte in storage at address EA +6 , and so forth until
- byte 7 of store_data is placed into the byte in storage at address EA.

\section*{Special Registers Altered None}

VSR Data Layout for stxsdx
src = VSR[XS]
\begin{tabular}{|l|l|}
\hline & unused \\
\hline 0 & 64 \\
\hline
\end{tabular}

Store VSX Scalar as Integer Byte Indexed
X-form

```

if SX=0 \& MSR.VSX=0 then VSX_Unavailable()
if SX=1 \& MSR.VEC=O then Vector_Unavailable()
EA\leftarrow(|RA=0) ? O:GPR[RA]) +GPR[RB]
MEM[EA,1) \leftarrowVSR[32xSX+S]. byte[7]

```

Let \(X S\) be the value \(32 \times 5 X+5\).
Let the effective address ( \(E A\) ) be sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

The contents of byte element 7 of VSR[ XS] are placed into the byte in storage addressed by EA.

\section*{Special Registers Altered:}

None
VSR Data Layout for stxsibx
sic = VSR[XS]
\begin{tabular}{|c|c|c|c|c|c|}
\hline & unused & byte & & unused & \\
\hline 0 & & \multicolumn{2}{|l|}{5664} & \multicolumn{2}{|r|}{127} \\
\hline
\end{tabular}

Store VSX Scalar as Integer Halfword Indexed X-form
Stxsihx
\begin{tabular}{|c|c|c|c|c|c|cc|c|}
\hline 31 & & S S,RA,RB \\
0 & & 6 & & RA & RB & & 941 & \begin{tabular}{l} 
SX \\
31
\end{tabular} \\
\hline
\end{tabular}
```

if $S X=0$ \& MSR.VSX=O then VSX_Unavailable(l)
if $S X=1$ \& MSR.VEC=O then Vector Unavailablell
$E A \leftarrow((R A=0) ? O: G P R[R A])+G P R[R B]$
MEM(EA, 2) $\leftarrow \operatorname{VSR}[32 \times 5 X+S]$.hword[3]

```

Let \(X S\) be the value \(32 \times 5 X+5\).
Let the effective address ( EA ) be sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

The contents of halfword element 3 of VSR[XS] are placed into the halfword in storage addressed by EA.

\section*{Special Registers Altered:}

None
VSR Data Layout for stxsihx
sIC = VSR[ XS]
\begin{tabular}{|l|c|c|c|}
\hline unused &, hword 3\(]\) & unused \\
\hline 0 & 48 & 64 & 127 \\
\hline
\end{tabular}

\section*{Store VSX Scalar as Integer Word Indexed X-form}
stxsiwx
\begin{tabular}{|c|c|c|c|c|cc|}
\hline 31 & 3 & S & RA, RB \\
0 & 31 & 6 & & RA & 16 & RB \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablell
\(E A \leftarrow((R A=0) ? 0: G P R[R A])+G P R[R B]\)
MEM \((E A, 4) \leftarrow \operatorname{VSR}[32 \times 5 X+5]\) word[1]

Let \(X S\) be the value \(32 \times 5 x+5\).
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

Let store_data be the contents of word element 1 of VSR[ XS].

When Big-Endian byte ordering is employed, store_data is placed in the word in storage at address EA in such order that;
- byte 0 of store_data is placed into the byte in storage at address EA,
- byte 1 of store_data is placed into the byte in storage at address \(E A+1\),
- byte 2 of store_data is placed into the byte in storage at address EA +2 , and
- byte 3 of store_data is placed into the byte in storage at address \(E A+3\).

When Little-Endian byte ordering is employed, store_data is placed in the word in storage at address EA in such order that;
- byte 0 of store_data is placed into the byte in storage at address \(E A+3\),
- byte 1 of store_data is placed into the byte in storage at address \(E A+2\),
- byte 2 of store_data is placed into the byte in storage at address EA +1 , and
- byte 3 of store_data is placed into the byte in storage at address EA.

Special Registers Altered None

VSR Data Layout for stxsiwx
src = VSR[ XS]
\begin{tabular}{|l|l|lr|}
\hline unused &. word [1] & unused \\
\hline 0 & 32 & 64 & 127 \\
\hline
\end{tabular}

Store VSX Scalar Single DS-form
\begin{tabular}{l} 
VRS, DS(RA) \\
\begin{tabular}{|c|c|c|ccc|c|}
\hline 61 & VRS & RA & & DS & \begin{tabular}{c}
3 \\
0
\end{tabular} & \({ }^{6}\)
\end{tabular} \\
\hline
\end{tabular}

> if MSR. VEC=0 then Vector_Unavailable()
> \(E A \leftarrow((\) RA=0)? \(0:\) GPR[RA] \()+\) EXTS(DS \(|\mid 0 b 00)\)
> MEM \((E A, 4) \leftarrow\) ConvertSP64toSP(VSR[VRS+32]. dword[0])

Let \(X S\) be the value VRS +32 .
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if \(R A=0\), and the signed integer value \(D S \| O b 00\).

Let store_data be the double-precision floating-point value in doubleword element 0 of VSR[ XS] converted to single-precision format

When Big-Endian byte ordering is employed, store_data is placed in the word in storage at address EA in such order that;
- byte 0 of store_data is placed into the byte in storage at address EA,
- byte 1 of store_data is placed into the byte in storage at address \(E A+1\),
- byte 2 of store_data is placed into the byte in storage at address EA +2 , and
- byte 3 of store_data is placed into the byte in storage at address \(E A+3\).

When Little-Endian byte ordering is employed, store_data is placed in the word in storage at address EA in such order that;
- byte 0 of store_data is placed into the byte in storage at address \(E A+3\),
- byte 1 of store_data is placed into the byte in storage at address \(E A+2\),
- byte 2 of store_data is placed into the byte in storage at address EA +1 , and
- byte 3 of store_data is placed into the byte in storage at address EA.

\section*{Special Registers Altered:} None

VSR Data Layout for stxssp
src = VSR[XS]


\section*{Store VSX Scalar Single-Precision Indexed X-form}
stxsspx
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & & S & RA & & RB & \\
\hline 0 & & 6 & & 11 & 16 & \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablel)
\(E A \leftarrow((R A=0) ? 0: G P R[R A])+G P R[R B]\)

MEM(EA, 4) \(\leftarrow\) ConvertSP64toSP(VSR[ \(32 \times 5 X+S]\). dword[0])

Let \(X S\) be the value \(32 \times 5 X+S\).
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

Let storedata be the double-precision floating-point value in doubleword element 0 of VSR[ XS] converted to single-precision format

When Big-Endian byte ordering is employed, store_data is placed in the word in storage at address EA in such order that;
- byte 0 of store_data is placed into the byte in storage at address EA,
- byte 1 of store_data is placed into the byte in storage at address \(E A+1\),
- byte 2 of store_data is placed into the byte in storage at address EA +2 , and
- byte 3 of store_data is placed into the byte in storage at address \(E A+3\).

When Little-Endian byte ordering is employed, store_data is placed in the word in storage at address EA in such order that;
- byte 0 of store_data is placed into the byte in storage at address \(E A+3\),
- byte 1 of store_data is placed into the byte in storage at address \(E A+2\),
- byte 2 of store_data is placed into the byte in storage at address EA +1 , and
- byte 3 of store_data is placed into the byte in storage at address EA.

\section*{Special Registers Altered} None

VSR Data Layout for stxsspx
src = VSR[ XS]
\begin{tabular}{|l|l|}
\hline & unused \\
\hline 0 & 64 \\
\hline
\end{tabular}

\section*{Store VSX Vector Byte*16 Indexed X-form}
stxvb16x \begin{tabular}{l} 
XS,RA,RB \\
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & 6 & S & RA & RB & & 1004 \\
0 & & & 11 & & 21 & \\
31
\end{tabular}\(|\)
\end{tabular}
if \(S X=0\) \& MSR. VSX=O then VSX_Unavailable ()
if \(S X=1\) \& MSR.VEC=O then Vector_Unavailablel)
\(E A \leftarrow R A=0\) ? GPR[ RB]: GPR[ RA] + GPR[ RB]
do \(\mathrm{i}=0\) to 15
\(\operatorname{MEM}(E A+i, 1) \leftarrow \operatorname{VSR}[32 \times S X+S]\), byte[i]
end

Let \(X S\) be the value \(32 \times 5 X+S\).
Let the effective address ( \(E A\) ) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

For each integer value from 0 to 15 , do the following.
The contents of byte element i of VSR[XS] are placed into the byte in storage at address EA +i .

\section*{Special Registers Altered:}

None

\section*{Example: Storing data using Store VSX Vector Byte*16 Indexed}
char X[16];
VSR[X]:


Storing a vector of 16 byte elements from VSR[ XS] into Big-Endian storage using sxvb16x, retaining left-to-right element ordering.
```


# Assumptions

# GPR[PX] = address of X

stxub16x xX,r0,rPX

```


Loading a vector of 16 byte elements from Little-Endian storage in VSR[XT] using Ixvb16x, retaining left-to-right element ordering.
\# Assumptions
\# GPR[PX] = address of \(X\)
stxub16x \(\quad x X, r 0, r P X\)
Little-endian storage image of \(X\)


\section*{Store VSX Vector Doubleword*2 Indexed \(X\)-form}
stxvd2x
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & & S & RA & RB & & 972 \\
\hline 0 & & 6 & & 11 & 16 & \\
\hline
\end{tabular}
\begin{tabular}{ll} 
a\{0:63\} & \(\leftarrow(R A=0) ? 0:\) GPR[RA] \\
EA\{0:63\} & \(\leftarrow a+G P R[R B]\) \\
MEM \((E A, 8)\) & \(\leftarrow\) VSR[32xSX+S].dword[0] \\
MEM \((E A+8,8)\) & \(\leftarrow\) VSR[32xSX+S].dword[1]
\end{tabular}

Let \(X S\) be the value \(32 \times 5 X+S\).
Let EA be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

For each integer value i from 0 to 1 , do the following.
Let store data be the contents of doubleword element i of VSR[ XS].

When Big-Endian byte ordering is employed, store_data is placed in the doubleword in storage at address EA \(88 \times i\) in such order that;
- byte 0 of store_data is placed into the byte in storage at address \(E A+8 \times i\),
- byte 1 of store_data is placed into the byte in storage at address EA \(+8 \times i+1\), and so forth until
- byte 7 of store_data is placed into the byte in storage at address EA \(+8 \times i+7\).

When Little-Endian byte ordering is employed, store_data is placed in the doubleword in storage at address EA \(88 \times i\) in such order that;
- byte 0 of store_data is placed into the byte in storage at address \(E A+8 \times i+7\),
- byte 1 of store_data is placed into the byte in storage at address EA \(+8 \times i+6\), and so forth until
- byte 7 of store_data is placed into the byte in storage at address EA \(+8 \times 1\).

\section*{Special Registers Altered}

None

VSR Data Layout for stxvd2x
src = VSR[XS]
\begin{tabular}{|l|l|}
\hline &.\(d\) word [0] \\
\hline 0 & 64 \\
\hline
\end{tabular}

Store VSX Vector Halfword*8 Indexed X-form
stxvh8x
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & 6 & S & RA & RB & & 940 \\
\hline 0 & & & 11 & & 16 & 21 \\
31
\end{tabular}
```

if SX=0 \& MSR.vSX=0 then VSX_Unavailable()
if SX=1 \& MSR.VEC=0 then Vector_Unavailable(l
EA\leftarrowRA=O?GPR[RB]:GPR[RA] + GPR[RB]
do i = 0 to 7
MEM[EA+2xi, 2) \leftarrowVSR[32xSX+5], hword[i]
end

```

Let \(X S\) be the value \(32 \times 5 X+S\).
Let the effective address (EA) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

For each integer value from 0 to 7 , do the following.
The contents of byte element i of VSR[XS] are placed into the byte in storage at address \(E A+i\).

For each integer value from 0 to 7, do the following. When Big-Endian byte ordering is employed, the contents of halfword element i of VSR[XS] are placed into the halfword in storage at address EA +2xi in such an order that;
- the contents of byte sub-element 0 of halfword element \(i\) of VSR[ XS] are placed into the byte in storage at address EA \(+2 \times \mathrm{i}\), and
- the contents of byte sub-element 1 of halfword element \(i\) of VSR[ XS] are placed into the byte in storage at address EA \(+2 \times i+1\).

When Little-Endian byte ordering is employed, the contents of halfword element i of VSR[XS] are placed into the halfword in storage at address EA \(+2 \times\) i in such an order that;
- the contents of byte sub-element 1 of halfword element \(i\) of VSR[ XS] are placed into the byte in storage at address EA+2xi, and
- the contents of byte sub-element 0 of halfword element \(i\) of VSR[ XS] are placed into the byte in storage at address EA \(+2 \times i+1\).

\section*{Special Registers Altered:}

None

\section*{Example: Storing data using Store VSX Vector Halfword*8 Indexed \\ short \(X[8]\);}

VSR[X]:


Storing a vector of 8 halfword elements from VSR[ X] into Big-Endian storage using stxvh8x, retaining left-to-right element ordering.
```


# Assumptions

# GPR[PX] = address of X

stxvh8x xX,ro,rpX

```
\[
\begin{aligned}
& \text { Big-endian storage image of } X
\end{aligned}
\]

Storing a vector of 8 halfword elements from VSR[ X] into Little-Endian storage using stxvh8x, retaining left-to-right element ordering.
```


# Assumptions

# GPR[PX] = address of X

stxvh8x XX,r0,rPX
Little-endian storage image of $X$

```



Store VSX Vector Word*4 Indexed X-form
stxvw4x XS,RA,RB


Let \(X S\) be the value \(32 \times 5 X+S\).
Let \(E A\) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

For each integer value i from 0 to 3 , do the following.
Let store data be the contents of word element
of VSR[ XS].
When Big-Endian byte ordering is employed, store_data is placed in the word in storage at address EA \(+4 \times \mathrm{i}\) in such order that;
- byte 0 of store_data is placed into the byte in storage at address EA \(+4 \times \mathrm{xi}\),
- byte 1 of store_data is placed into the byte in storage at address EA \(+4 \times i+1\), and so forth until
- byte 3 of store_data is placed into the byte in storage at address EA \(+4 \times i+3\).

When Little-Endian byte ordering is employed, store_data is placed in the word in storage at address \(E A+4 \times i\) in such order that;
- byte 0 of store_data is placed into the byte in storage at address EA \(+4 \times i+3\),
- byte 1 of store_data is placed into the byte in storage at address EA \(+4 \times i+2\), and so forth until
- byte 3 of store_data is placed into the byte in storage at address \(E A+4 \times i\).

\section*{Special Registers Altered \\ None}

\section*{VSR Data Layout for stxvw4x}

SIC = VSR[XS]
\begin{tabular}{|l|l|l|l|}
\hline. word[ 0] & .word[ 1] &. word[ 2] & .word[ 3] \\
\hline 0 & 32 & 64 & 96
\end{tabular}
Extended Mnemonic \(\quad\) Equivalent To
stxvx XS,RA,RB stxvd2x XS,RA, RB

Usage: stxvx can be used for vector store operations when using Big-Endian byte-ordering, independent of element size

\section*{Store VSX Vector DQ-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{stxv XS, DQ(RA)} \\
\hline \[
01
\] & \[
{ }_{6} \mathrm{~S}
\] & \({ }_{11} \mathrm{RA}\) & 16 & DQ & \begin{tabular}{|c|c|}
\hline\(X\) & 5 \\
28 & 29 \\
\hline 18 \\
\hline
\end{tabular} \\
\hline
\end{tabular}
if \(S X=0\) \& MSR.VSX \(=0\) then VSX_Unavailablell
if \(S X=1\) \& MSR.VEC=O then Vector_Unavailable(l)
\(E A \leftarrow R A=0\) ? GPR[RB]: GPR[RA] +EXTS(DQ||ObOOOO)
\(\operatorname{MEM}(E A, 16) \leftarrow \operatorname{VSR}[32 \times 5 X+5]\)

Let \(X S\) be the value \(32 \times 5 X+S\).
Let EA be the sum of the contents of GPR[RA], or 0 if \(R A=0\), and the signed integer value \(D Q \ll 4\).

Letstore_data be the contents of VSR[ XS].
When Big-Endian byte ordering is employed, store_data is placed into the quadword in storage at address EA in such an order that;
- byte 0 of store_data is placed into the byte in storage at address EA,
- byte 1 of store_data is placed into the byte in storage at address \(E A+1\), and so forth until
- byte 15 of store_data is placed into the byte in storage at address \(E A+15\).

When Little-Endian byte ordering is employed, store_data is placed into the quadword in storage at address \(E A\) in such an order that;
- byte 15 of store_data is placed into the byte in storage at address EA,
- byte 14 of store_data is placed into the byte in storage at address EA +1 , and so forth until
- byte 0 of store_data is placed into the byte in storage at address EA +15 .

\section*{Special Registers Altered None}

\section*{Store VSX Vector with Length X-form}

```

if SX=0 \& MSR.VSX=0 then VSX_Unavailable()
if SX=1 \& MSR.VEC=O then Vector_Unavailable()
EA\leftarrow(RA=0) ? 0: GPR[RA]
nb \leftarrow EXTZ(GPR[ RB], bit[0:7])
if MSR.LE = 0 then |/ Big-Endian byte-ordering
store_data \& VSR[32xSX+S], byte[0:nb-1]
else II Little-Endian byte ordering
store_data }\leftarrow\mathrm{ VSR[ 32xSX+S], byte[16.nb:15]
MEM(EA, nb) \& store_data

```

Let \(X S\) be the value \(32 \times S X+S\).
Let the effective address (EA) be the contents of \(G P R[R A]\), or 0 if \(R A\) is equal to 0 .

Let \(n b\) be the unsigned integer value in bits \(0: 7\) of GPR[RB].

If \(n b\) is equal to 0 , the storage access is not performed.
Otherwise, when Big-Endian byte-ordering is employed, do the following.

If nb less than 16 , the contents of the leftmost nb bytes of VSR[ XS] are placed in storage starting at address EA.

Otherwise, the contents of VSR[ XS] are placed into the quadword in storage at address EA.

Otherwise, when Little-Endian byte ordering is employed, do the following.

If nb less than 16 , the contents of the rightmost nb bytes of VSR[ XS] are placed in storage starting at address EA in byte-reversed order.

Otherwise, the contents of VSR[XS] are placed into the quadword in storage at address EA in byte-reversed order.

If the contents of bits \(8: 63\) of GPR[RB] are not equal to 0 , the results are boundedly undefined.

\section*{Special Registers Altered: \\ None}

Example: Storing less than 16-byte data from VSR using stxvl
```

char S[14] = "This is a TEST";
short X[6] = { OXEOE1, OXE2E3, OXE4E5, OXE6E7, OXE8E9, OXEAEB };
binary80 Z = OXFOF1F2F3F4F5FGF7F8F9

```

Storing less than 16-byte data in VSR[XS] into Big-Endian storage using stxvl.
```


# Assumptions

# GPR[NS] = 14 (length of S in \# of bytes)

# GPR[NX] = 12 (Iength of X in \# of bytes)

# GPR[NZ] = 10 (length of Z in \# of bytes)

# GPR[PS] = address of S

```

VSR register image of \(\mathrm{S}, \mathrm{X}\), \& \(Z\)

VSR[S]
VSR[X]:
VSR[ Z]:
F0 F1 F2 F3 F4 F5 F6 F7 F8 F9 000000000000

add \(\quad\) PPX,rPS,rNS \# address of \(X\)
add rPZ,rPX,rNX \# address of \(Z\)
sldi rLS,rNS,56
sldi rLX,rNX,56
sldi rLZ,rNZ,56
stxul XS,rPS,rLS
stxul XX,rPX,rLX
stxvl XZ,rPZ,rLZ
Final state of Big-Endian storage image of \(\mathrm{S}, \mathrm{X}, \& \mathrm{Z}\)
\(\operatorname{addr}(S)+0 \times 0000\) :

addr (S) \(+0 \times 0020\)

Storing less than 16-byte data in VSR[XS] into Little-Endian storage using stxvl.


Final state of Little-Endian storage image of \(\mathrm{S}, \mathrm{X}, \& \mathrm{Z}\)

\(\operatorname{addr}(S)+0 \times 0010:\)
\(\operatorname{addr}(S)+0 \times 0020:\)


Store VSX Vector Left-justified with Length X-form
Stxvll
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & & S S RA,RB \\
\hline 0 & & 6 & & RA & & RB \\
\hline 11 & & & 429 & SX \\
31 \\
\hline
\end{tabular}
if \(S X=0\) \& MSR.VSX=O then VSX_Unavailable (l)
if \(S X=1\) \& MSR.VEC=O then Vector_Unavailablell
\(E A \leftarrow(R A=0) ? O: G P R[R A]\)
\(n b \leftarrow \operatorname{EXTZ}(\) GPR \([R B]\), bit \([0: 7])\)
If nb>0 then do \(i=0\) to nb. 1
MEM \((E A+i, 1) \leftarrow \operatorname{VSR}[32 \times S X+S]\), byte[i]
end
Let \(X S\) be the value \(32 \times 5 X+S\).
Let the effective address (EA) be the contents of \(G P R[R A]\), or 0 if \(R A\) is equal to 0 .

Let \(n b\) be the unsigned integer value in bits \(0: 7\) of GPR[RB].

If \(n b\) is equal to 0 , the storage access is not performed.
Otherwise, do the following.
If \(n b\) less than 16 , the contents of the leftmost \(n b\) bytes of VSR[ XS] are placed in storage starting at address EA.

Otherwise, the contents of VSR[ XS] are placed into the quadword in storage at address EA.

Data is stored from VSR[XS] into storage in Big-Endian byte ordering (i.e., the contents of byte element 0 of VSR[ XS] are placed into the byte in storage at address EA, the contents of byte element 1 of VSR[ XS] are placed into the byte in storage at address \(E A+1\), and so forth).

If the contents of bits \(8: 63\) of GPR[RB] are not equal to 0 , the results are boundedly undefined.

\section*{Special Registers Altered:}

None

\section*{Example: Storing less than 16-byte left-justified data}
\(\begin{array}{ll}\text { decimal } & X=+1234567890123456789 ; \\ \text { decimal } & Y=-123456 ; \\ \text { decimal } & Z=+1004966723510220 ;\end{array}\)
Storing less than 16-byte data, left-justified in VSR[ XS], into storage using stxvll.
\# Assumptions
\# GPR[NX] \(=10\) (Iength of \(X\) )
\# GPR[NY] \(=4\) (length of \(Y\) )
\# GPR[NZ] \(=9\) (length of \(Z\) )
\# GPR[PX] = address of \(X\)
\# GPR[PY] \(=\) address of \(Y=\) address of \(X+10\)
\# GPR[PZ] = address of \(Z=\) address of \(X+10+4\)

VSR[X]: 0134677890123456789 C 000000000000
VSR[Y]: 0123456 D 000000000000000000000000
VSR[2]:


Initial state of Big-endian \& Little-Endian storage image of \(X, Y, \& Z\)

stxull \(\quad \mathrm{XX}, \mathrm{rPX}, \mathrm{rNX}\)
stxull \(x Y, r P Y, r N Y\)
stxull \(x Z, r P Z, r N Z\)
Final state of Big-endian \& Little-Endian storage image of \(X, Y, \& Z\)


\section*{Store VSX Vector Indexed X-form}
stxvx XS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|}
\hline & 31 & & S & & RA & R & & & 396 & sx \\
\hline 0 & & 6 & & 11 & & & & 1 & & 31 \\
\hline
\end{tabular}
if \(S X=0\) \& MSR. VSX=0 then VSX_Unavailable el)
if \(S X=1\) \& MSR.VEC=O then Vector_Unavailable el)
\(E A \leftarrow R A=0\) ? GPR[ \([B]\) : GPR[RA] + GPR[ RB]
MEM \(\operatorname{EA}, 16) \leftarrow \operatorname{VSR}[32 \times S X+5]\)

Let \(X S\) be the value \(32 \times 5 x+S\).
Let the effective address ( \(E A\) ) be the sum of the contents of GPR[RA], or 0 if RA is equal to 0 , and the contents of GPR[RB].

When Big-Endian byte ordering is employed, store_data is placed into the quadword in storage at address EA in such an order that;
- byte 0 of store_data is placed into the byte in storage at address EA,
- byte 1 of store_data is placed into the byte in storage at address EA+1, and so forth until
- byte 15 of storedata is placed into the byte in storage at address EA+15.

When Little-Endian byte ordering is employed, store_data is placed into the quadword in storage at address EA in such an order that;
- byte 15 of storedata is placed into the byte in storage at address EA,
- byte 14 of store_data is placed into the byte in storage at address \(E A+1\), and so forth until
- byte 0 of store_data is placed into the byte in storage at address EA +15 .

Special Registers Altered:
None

\section*{Version 3.0}

\section*{Example: Storing data using Store VSX Vector Indexed}
```

char W[16] = { OXFO, OXF1, OXF2, OXF3, OXF4, OXF5, OXF6, OXF7, OXEO, OXE1, OXE2, OxE3, OxE4, OXE5, OXE6, OxE7 };
short X[8]={OXFOF1, OXF2F3, OXF4F5, OXF6F7, OXEOE1, OXE2E3, OXE4E5,OXE6E7 };
float Y[4] = { OXFOF1_F2F3, OXF4F5_F6F7, OXEOE1_E2E3, OXE4E5_E6E7 };
double Z[2] = {OXFOF1_F2F3_F4F5_F6F7,OXEOE1_E2E3_E4E5_E6E7 };

```

Storing 16 bytes of data into Big-Endian storage from VSR[ XS] using stxvx


\section*{\# Assumptions}
\# GPR[PW] = address of W
\# GPR[PX] = address of \(X=\) GPR[PW] +16
\# GPR[PY] = address of \(Y=\) GPR[PW] +32
\# GPR[PZ] \(=\) address of \(Z=\) GPR[PW] +48
stevx xW,ro,rPW
stevx \(\quad X X, r 0, r P X\)
stevx \(\quad X Y, r 0, r P Y\)
stxvx \(X Z, r 0, r P Z\)
Big-endian storage image of \(\mathrm{W}, \mathrm{X}, \mathrm{Y}, \& \mathrm{Z}\)
addr (W+0×0000):
addr (W+0x0010):
addr (W+0x0020):
addr (W+0x0030):

Storing 16 bytes of data into Little-Endian storage from VSR[ XS] using stxvx.


VSR[Y]:
E4 E5 E6 E7 E0 E1 E2 E3 F4 F5 F6 F7 F0 F1 F2 F3
VSR[z]:
E0 E1 E2 E3 E4 E5 E6 E7 F0 F1 F2 F3 F4 F5 F6 F7

\# Assumptions
\# GPR[PW] = address of \(W\)
\# GPR[PX] \(=\) address of \(X=\) GPR[PW] +16
\# GPR[PY] = address of \(Y=\) GPR[PW] +32
\# GPR[PZ] = address of \(Z=\) GPR[PW] +48
stavx \(\quad \mathrm{xW}, \mathrm{ro}, \mathrm{rPW}\)
stevx \(\quad x, r o, r \operatorname{PX}\)
stxux XY,ro,rPY
stxux \(x z, r 0, r p z\)
addr (W+0x0000):
addr (W+0x0010):
addr (W+0x0020):
addr (W+0x0030):

Little-endian storage image of \(\mathrm{W}, \mathrm{X}, \mathrm{Y}, \& \mathrm{Z}\)


VSX Scalar Absolute Value Double-Precision XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{xsabsdp \(\quad\) XT, XB} \\
\hline 060 & \(6^{\text {T }}\) & \[
{ }_{11} \quad \text { III }
\] & 16 & B & 21 & 345 & \(|\)\begin{tabular}{l} 
BX \\
\(30 \times\) \\
\\
\hline 1
\end{tabular} \\
\hline
\end{tabular}
if MSR.VSX=O then VSX_Unavailablell
result.bit[0] \(\leftarrow\) ObO
result. bit \([1: 63] \leftarrow\) VSR[ \(32 \times B X+B]\), dwor d[0], bit [1:63]
VSR[32xTX + T]. dword[ 0\(] \leftarrow\) result
VSR[32xTX+T].dwor d[1] © OxUUUU_UUUU_UUUU_UUUU
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
The absolute value of the double-precision floating-point operand in doubleword element 0 of VSR[ XB] is placed into doubleword element 0 of VSR[ XT] in double-precision format.

The contents of doubleword element 1 of VSR[ XT] are undefined.

Special Registers Altered
None

VSR Data Layout for xsabsdp
\(\operatorname{src}=\mathrm{VSR}[X B]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\) VSR[ \(X T]\)


\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.
\begin{tabular}{|l|ll|}
\hline DP & undefined \\
\hline 0 & 64 & 127 \\
\hline
\end{tabular}

VSX Scalar Absolute Quad-Precision X-form
xsabsqp VRT,VRB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline 063 & \[
{ }_{6} \text { VRT }
\] & 11 & 0 & 16 & VRB & 21 & 804 & / \\
\hline
\end{tabular}
if MSR.VSX=O then VSX_Unavailable(l)

VSR[VRT+32] \(\leftarrow\) VSR[VRB 32\(]\) \& Ox7FFF_FFFF_FFF_FFFF_FFFF_FFF_FFFFFFFF
Let \(X T\) be the value VRT +32 .
Let \(X B\) be the value VRB +32 .
The absolute value of the quad-precision floating-point value in VSR[ XB] is placed into VSR[ XT].

\section*{Special Registers Altered:}

None
VSR Data Layout for xsabsqp
VSR[ XB]
sic

VSR[ XT]
tgt
\(\square\) tgt

VSX Scalar Add Double-Precision XX3-form
xsadddp XT,XA,XB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline -60 & 6 & T & 11 & A & 16 & B & 21 & 32 &  \\
\hline
\end{tabular}

\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{if( ~vex_flag ) then do} \\
\hline & ↔result || 0xUUUU_UUUU_UUUU_UUUU \\
\hline FPR & \(\leftarrow\) ClassSP(result) \\
\hline FR & \(\leftarrow\) inc_flag \\
\hline FI & \(\leftarrow\) xx_flag \\
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{\[
\begin{aligned}
& \text { end } \\
& \text { else do }
\end{aligned}
\]}} \\
\hline & \\
\hline FR & \(\leftarrow 0 \mathrm{~b} 0\) \\
\hline FI & \(\leftarrow 0 \mathrm{bo}\) \\
\hline end & \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src2 is added \({ }^{[1]}\) to src1, producing a sum having unbounded range and precision.

The sum is normalized \({ }^{[2]}\).
See Table 57, "Actions for xsadddp," on page 515.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

\footnotetext{
1. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( \(G, R\), and \(X\) ) enter into the computation.
2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \(v \leftarrow-\) Infinity & \(v \leftarrow-\) Infinity & \(v \leftarrow\) - Infinity & \(\mathrm{V} \leftarrow\)-Infinity & \(v \leftarrow-\) Infinity & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vxisi_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
\mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{array}
\] \\
\hline -NZF & \(v \leftarrow-\) Infinity & \(v \leftarrow A(\operatorname{src} 1, \mathrm{src} 2)\) & \(\mathrm{v} \leftarrow \mathrm{SrC1}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{A}(\mathrm{src} 1, \mathrm{src} 2)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{src} 2\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { srcc } 2) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline -Zero & \(v \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \(v \leftarrow\)-Zero & \(v \leftarrow\) Rezd & \(\mathrm{v} \leftarrow \mathrm{SrC2}\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src2 }) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\overline{\text { - }}\) +Zero & \(v \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \(v \leftarrow\) Rezd & \(v \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src2 }) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline ¢ +NZF & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{A}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{A}(\mathrm{src} 1, \mathrm{src} 2)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q} \text { (src2) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline +Infinity & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{dQNaN} \\
& \text { vxisi_flag } \leftarrow 1
\end{aligned}
\] & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline QNaN & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \[
\begin{array}{|l|l}
\hline \mathrm{V} \leftarrow \text { src1 } \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline SNaN & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow \mathrm{Q} \text { (src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{array}{|l|}
\hline v \leftarrow Q(s r c 1) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \[
\begin{array}{|l|}
\hline v \leftarrow Q(\text { src1 } 1) \\
v x \text { nnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src1 } 1) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src1 } 1) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \[
\begin{array}{|l|}
\hline v \leftarrow Q(\text { src1 } 1) \\
v x s n a n \_f l a g ~
\end{array} 1
\] & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
\mathrm{vxsnan} \text { _flag } \leftarrow 1 \\
\hline
\end{array}
\] & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline Explanation: & & & & & & & & \\
\hline sicl & \multicolumn{8}{|l|}{The double-precision floating-point value in doubleword element 0 of VSR[ XA] .} \\
\hline sic2 & \multicolumn{8}{|l|}{The double-precision floating-point value in doubleword element 0 of VSR[ XB] .} \\
\hline dQNaN & \multicolumn{8}{|l|}{Default quiet \(\mathrm{NaN}\left(0 \times 7 \mathrm{FF} 8.0000 \_0000 \_0000\right)\).} \\
\hline NZF & \multicolumn{8}{|l|}{Nonzero finite number.} \\
\hline Read & \multicolumn{8}{|l|}{Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs).} \\
\hline \(A(x, y)\) & \multicolumn{8}{|l|}{Return the normalized sum of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. Note: If \(x=\cdot y, v\) is considered to be an exact-zero-difference result (Rezd).} \\
\hline \(Q(x)\) & \multicolumn{8}{|l|}{Return a QNaN with the payload of \(x\).} \\
\hline \(\checkmark\) & \multicolumn{8}{|l|}{The intermediate result having unbounded signficand precision and unbounded exponent range.} \\
\hline
\end{tabular}

Table 57.Actions for xsadddp

Version 3.0
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multirow[b]{2}{*}{Range of \(\mathbf{v}\)} & \multirow[b]{2}{*}{Case} & \multicolumn{5}{|c|}{Rounding Mode} \\
\hline & & Round To Nearest (RTN) & \[
\begin{aligned}
& \text { Round Towards } \\
& \text { Zero (RTZ) }
\end{aligned}
\] & Round Towards +Infinity (RTP) & Round Towards -Infinity (RTM) & Round To Odd (RTO) \\
\hline \(v\) is Q QNaN & Special & r + V & \(r\) + 1 & \(1 \leqslant v\) & \(t \leqslant v\) & r + 1 \\
\hline \(v=. \mid n f i n i t y\) & Special & \(1+1\) & \(r+1\) & \(1+1\) & \(t \leftarrow v\) & \(1+1\) \\
\hline - Infinity < V \(\leq\) - ( \(\|\) max +1 ulp \()\) & Overflow & \[
\begin{aligned}
& \hline 9 \leftarrow \operatorname{nd}(v) \\
& i \leftarrow \cdot \operatorname{lnfinity} \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& \hline 9 \leqslant \operatorname{rid}(v) \\
& 1 \leftarrow \cdot \max \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& 9 \leftarrow \mathrm{rnd}(v) \\
& 1 \leftarrow \cdot N \max \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& 9 \leftarrow \operatorname{rnd}(v) \\
& 1 \leftarrow \cdot \operatorname{lnfinity} \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& 9 \leftarrow \operatorname{rnd}(v) \\
& r \leftarrow \cdot \operatorname{Nmax} \\
& \hline
\end{aligned}
\] \\
\hline \multirow[t]{2}{*}{} & Overflow & \[
\begin{aligned}
& \hline 9 \leftarrow \operatorname{lnd}(v) \\
& i \leftarrow \cdot \operatorname{lnf} \mathrm{inity}
\end{aligned}
\] & - & - & \[
\begin{aligned}
& 9 \leftarrow \operatorname{rnd}(v) \\
& 1 \leftarrow \cdot \operatorname{lnf} \mathrm{inity}
\end{aligned}
\] & \\
\hline & Normal & - & \(1 \leftarrow \cdot \mathrm{Mmax}\) & \(1 \leftarrow \cdot \mathrm{Mmax}\) & - & \(\mathrm{r} \leftarrow \cdot \mathrm{Mmax}\) \\
\hline \multirow[t]{2}{*}{} & Overflow & - & - & - & \[
\begin{aligned}
& 9 \leqslant \operatorname{lnd}(v) \\
& 1 \leftarrow \cdot \operatorname{lnfinity} \\
& \hline
\end{aligned}
\] & \\
\hline & Normal &  & \(\underline{r}\) + \(\max ^{\text {a }}\) & \(1+\cdot N\) max & - & \(\underline{1+\cdot \mathrm{Max}}\) \\
\hline \(\cdot \mathrm{Mmax} \leq \mathrm{V} \leq \cdot \mathrm{Vmin}\) & Normal & \(1 \leqslant \operatorname{lnd}(\mathrm{v})\) & \(1 \leqslant \operatorname{ind}(\mathrm{v})\) & \(1 \leftarrow \operatorname{ind}(\mathrm{v})\) & \(1 \leqslant \operatorname{ind}(\mathrm{v})\) & \(1 \leqslant \operatorname{ind}(\mathrm{v})\) \\
\hline - Nmin \(<\mathrm{V}\) < - Zero & Tiny &  & \[
\begin{aligned}
& 9 \leftarrow \operatorname{rnd}(v) \\
& i \leftarrow \operatorname{rod}(\operatorname{den}(v)) \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& 9 \leftarrow \operatorname{rnd}(v) \\
& 1 \leftarrow \operatorname{rnd}(\operatorname{den}(v)) \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& 9+\operatorname{rnd}(v) \\
& i \leftarrow \operatorname{rnd}(\operatorname{den}(v)) \\
& \hline
\end{aligned}
\] &  \\
\hline \(v=\cdot\) ero & Special & r + 1 & \(t \leqslant v\) & \(1 \leqslant v\) & \(t \leftarrow v\) & \(r\) i \(V\) \\
\hline \(v=\) Rezd & Special & \(1 \leqslant+\) erro & \(1 \leqslant+\) Eero & \(1 \leqslant+\) ero & \(1 \leftarrow\) Lero &  \\
\hline \(v=+2 e r o\) & Special & i + 1 & \(r\) + 1 & \(1 \leftarrow v\) & \(1 \leftarrow v\) &  \\
\hline +Zero < v < +Mmin & Tiny &  &  & \[
\begin{aligned}
& \hline 9 \leftarrow \operatorname{rnd}(v) \\
& i \leftarrow \operatorname{rnd}(\operatorname{den}(v)) \\
& \hline
\end{aligned}
\] & \[
\begin{array}{|l|l|}
\hline 9+\operatorname{rnd}(v) \\
1 & \operatorname{rid}(\operatorname{den}(v)) \\
\hline
\end{array}
\] & \[
\begin{aligned}
& \hline 9 \leftarrow \operatorname{rnd}(v) \\
& i \leftarrow \bmod (\operatorname{den}(v)) \\
& \hline
\end{aligned}
\] \\
\hline + Mmin \(\leq\) v \(\leq+4\) max & Normal & \(1 \leqslant \operatorname{lnd}(\mathrm{v})\) & \(r \leqslant \operatorname{lnd}(\mathrm{v})\) & \(1 \leftarrow \operatorname{lnd}(v)\) & \(1 \leqslant \operatorname{ind}(\mathrm{v})\) & \(1 \leqslant \operatorname{lnd}(\mathrm{v})\) \\
\hline \multirow[t]{2}{*}{} & Overflow & - & - & \[
\begin{aligned}
& 9 \leftarrow \operatorname{rnd}(v) \\
& i \leftarrow+\operatorname{lnfinity}
\end{aligned}
\] & - & \\
\hline & Normal & t + + 4 max & \(1++4\) max & - & \(1++1\) max & r + + 1 max \\
\hline \multirow[t]{2}{*}{} & Overflow & \[
\begin{aligned}
& \hline 9 \leftarrow \operatorname{nd}(v) \\
& 1 \leftarrow+\operatorname{nfinity} \\
& \hline
\end{aligned}
\] & - & \[
\begin{aligned}
& 9 \leftarrow \operatorname{rnd}(v) \\
& i \leftarrow+\operatorname{lnfinity} \\
& \hline
\end{aligned}
\] & - & \\
\hline & Normal & - & \(1 \leftarrow+4\) max & - & \(1 \leftarrow+4\) max & \(\underline{r}+4 \mathrm{mma}\) \\
\hline  & Overflow & \[
\begin{aligned}
& \hline 9 \leftarrow \operatorname{rd}(v) \\
& i \leftarrow+\operatorname{lnf} \text { inity }
\end{aligned}
\] & \[
\begin{aligned}
& 9 \leftarrow \operatorname{lnd}(v) \\
& 1 \leftarrow+\| \max
\end{aligned}
\] & \[
\begin{aligned}
& 9 \leftarrow \operatorname{rnd}(v) \\
& i \leftarrow+\operatorname{lnf} i n i t y
\end{aligned}
\] & \[
\begin{aligned}
& 9 \leftarrow \operatorname{rnd}(v) \\
& 1 \leftarrow+\| \max
\end{aligned}
\] & \[
\begin{aligned}
& 9 \leftarrow \text { rnd }(v) \\
& 1 \leftarrow+\| \max x
\end{aligned}
\] \\
\hline \(v=+\) nfinity & Special & \(i \leqslant v\) & r \(\leftarrow v\) & \(1 \leftarrow v\) & \(1 \leftarrow v\) & i \(\leftarrow v\) \\
\hline
\end{tabular}

\section*{Explanation:}

\section*{- This situation cannot occur.}
\(v \quad\) The precise intermediate result defined in the instruction having unbounded range and precision.
den \((x) \quad\) The significand of \(x\) is shifted right by the amount of the difference between the target rounding precision Emin and the unbiased exponent of \(x\). The unbiased exponent of the denormalized value is Emin. The significand of the denormalized value has unbounded significand precision.
\[
\begin{array}{ll}
E_{\min n}=-16382 & \text { (quad-precision) } \\
E_{\min n}=-16382 & \text { (double-extended-precision) } \\
E_{\min n}=-1022 & \text { (double-precision) } \\
\text { mi } n=-126 & \text { (single-precision) }
\end{array}
\]

Rezd Exact-zero-difference result. Applies only to add operations involving source operands having the same magnitude and different signs or subtract operations involving source operands having the same magnitude and same signs. Whether +Zero or -Zero is returned is controlled by the setting of the rounding mode in RN, even when the rounding mode is overridden to Round to Odd.
\(r n d(x) \quad\) The significand of \(x\) is rounded to the target rounding precision according to the rounding mode specified in FPSCR. RN. Exponent range of the rounded result is unbounded. See Section 7.3.2.6.
\(N \max \quad\) Largest (in magnitude) representable normalized number in the target rounding precision format.
\begin{tabular}{|c|c|c|}
\hline \(N\) max \(=\) & \(\times 1 . F F F F F F F F F F F F F F F F F F F F F F F F F F F F\) & ( \\
\hline \(N\) max \(= \pm 2^{+16383}\) & x 1.FFFFFFFFFFFFFFFO000000000000 & (double-extended-p \\
\hline \(N\) max \(= \pm 2^{+1023}\) & x 1.FFFFFFFFFFFFF000000000000000 & (double-precision) \\
\hline \(N\) max \(= \pm 2^{+127}\) & x 1.FFFFFF000000000000000000 & (single-precision) \\
\hline
\end{tabular}

Nmin Smallest (in magnitude) representable normalized number in the target rounding precision format.
Nmin \(= \pm 2.16382 \times 1.0000000000000000000000000000\)
Nmin \(= \pm 2.16382 \times 1.0000000000000000000000000000\)
Nmin \(= \pm 2.1022 \times 1.0000000000000000000000000000\)
(doud-precision)
Nmin \(n= \pm 2.126 \times 1.000000000000000000000000\)
(dingle-prendeded-precision)
ulp
Least significant bit in the target precision format's significand (Unit in the Last Position).

\section*{Table 58.Scalar Floating-Point Intermediate Result Handling}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline Case &  & \[
\begin{aligned}
& \ddot{0} \\
& \dot{\tilde{U}} \\
& \text { n } \\
& \frac{a}{L}
\end{aligned}
\] & \[
\begin{aligned}
& \text { 山 } \\
& \dot{\tilde{y}} \\
& \text { n } \\
& \frac{a}{L}
\end{aligned}
\] & \[
\begin{aligned}
& \dot{N} \\
& \dot{\tilde{U}} \\
& \text { n } \\
& \frac{n}{u}
\end{aligned}
\] & \[
\begin{aligned}
& \underset{x}{x} \\
& \dot{\tilde{U}} \\
& \text { n } \\
& \frac{n}{4}
\end{aligned}
\] & \[
\begin{aligned}
& 0 \\
& \frac{0}{\sigma} \\
& \frac{1}{4} \\
& \varepsilon^{\prime} \\
& \frac{\sigma}{n} \\
& \dot{x} \\
& >
\end{aligned}
\] & \[
\begin{aligned}
& 0 \\
& \frac{0}{4} \\
& \overline{4} \\
& \underline{N} \\
& \bar{x} \\
& >
\end{aligned}
\] & \[
\begin{aligned}
& 0 \\
& \frac{0}{\pi} \\
& \overline{4} \\
& \bar{n} \\
& \bar{x} \\
& >
\end{aligned}
\] & \[
\begin{aligned}
& 0 \\
& \frac{0}{4} \\
& \bar{\prime} \\
& \bar{\sigma} \\
& \bar{x} \\
& >
\end{aligned}
\] & \[
\begin{aligned}
& 0 \\
& 0 \\
& \frac{0}{4} \\
& 1 \\
& 0 \\
& 0 \\
& X \\
& >
\end{aligned}
\] & \[
\begin{aligned}
& 0 \\
& 0 \\
& \frac{0}{4} \\
& 1 \\
& \frac{1}{\sigma} \\
& 0 \\
& x \\
& >
\end{aligned}
\] & \[
\begin{aligned}
& \text { o } \\
& \frac{0}{4} \\
& \mathbf{x} \\
& \times
\end{aligned}
\] & \[
\text { Is } r \text { inexact? }(r \neq v)
\] &  & \[
\text { Is } q \text { inexact? }(q \neq v)
\] &  & Returned Results and Status Setting \\
\hline \multirow{16}{*}{Special} & - & - & - & 0 & - & - & - & - & - & - & - & 1 & - & - & - & - & T(r), class_bfp(r), fi(0), fi(0), fx \(2 \times\) ) \\
\hline & - & - & - & 1 & - & - & - & - & - & - & - & 1 & - & - & - & - & fx(zx), error (1) \\
\hline & 0 & - & - & - & - & - & - & - & - & - & 1 & - & - & - & - & - & T(r), class_bfp(r), fi(0), fr(0), fx(VXSQRT) \\
\hline & 0 & - & - & - & - & - & - & - & - & 1 & - & - & - & - & - & - & \(\mathrm{T}(\mathrm{r})\), class_bfp(r), fi(0), fr(0), fx fl VXZDZ) \\
\hline & 0 & - & - & - & - & - & - & - & 1 & - & - & - & - & - & - & - & T(r), class_bfplr), fi(0), fr(0), fx(VXIDI) \\
\hline & 0 & - & - & - & - & - & - & 1 & - & - & - & - & - & - & - & - & T(r), class_bfp(r), fi(0), fr(0), fx(VXISI) \\
\hline & 0 & - & - & - & - & 0 & 1 & - & - & - & - & - & - & - & - & - & \(\mathrm{T}(\mathrm{r})\), class_bfp(r), fi (0), fr(0), fx(VXIMZ) \\
\hline & 0 & - & - & - & - & 1 & 0 & - & - & - & - & - & - & - & - & - & T(r), class_bfp(r), fi(0), fr(0), fx(VXSNAN) \\
\hline & 0 & - & - & - & - & 1 & 0 & - & - & - & - & - & - & - & - & - & T(r), class_bfp(r), fi(0), fr(0), fx(VXSNAN), fx(VXIMZ) \\
\hline & 1 & - & - & - & - & - & - & - & - & - & 1 & - & - & - & - & - & fx(VXSQRT), error () \\
\hline & 1 & - & - & - & - & - & - & - & - & 1 & - & - & - & - & - & - & fx(VXZDZ), error () \\
\hline & 1 & - & - & - & - & - & - & - & 1 & - & - & - & - & - & - & - & fx(VX) Di), error () \\
\hline & 1 & - & - & - & - & - & - & 1 & - & - & - & - & - & - & - & - & fx(VX|SI), error () \\
\hline & 1 & - & - & - & - & 0 & 1 & - & - & - & - & - & - & - & - & - & fx(VXI MZ), error () \\
\hline & 1 & - & - & - & - & 1 & 0 & - & - & - & - & - & - & - & - & - & fx(VXSNAN), error() \\
\hline & 1 & - & - & - & - & 1 & 1 & - & - & - & - & - & - & - & - & - & fx(VXSNAN), fx(VXIMZ), error () \\
\hline
\end{tabular}

Explanation:
-
The results do not depend on this condition.
\(\mathrm{T}(\mathrm{x}) \quad\) Places the result into the target VSR.
For scalar single-precision and double-precision results
VSR[XT]. dword[0] \&bfp_CONVERT_TO_BFP64(r)
VSR[XT].dword[1] \& OxUUUU UUUU־UUÜU UUUU
For scalar quad-precision results
VSR[VRT+32] \& bf p_CONVERT_TO_BFP128(r)
class_bf \(p(x)\) Sets FPSCR. FPRF to the sign and class of \(x\).
FPSCR.FPRF \& fprf_CLASS_BFP32(x) (single-precision)
FPSCR.FPRF \(\leftarrow\) f Prf \(^{-}\)CLASS \({ }^{-}\)BFP64(x) (double-precision)
FPSCR.FPRF \& fprf_CLASS_BFP128(x) (quad-precision)
\(f x(x) \quad F P S C R . F X\) is set to 1 if \(F P S C R, x=0\). \(F P S C R . x\) is set to 1 .
\(f i(x) \quad\) FPSCR. \(F I\) is set to the value \(x\).
\(f r(x) \quad\) FPSCR. \(F R\) is set to the value \(x\).
\(\beta \quad\) Wrap adjust
\(\beta=2^{192}\) (single-precision)
\(\beta=2^{1536}\) (double-precision)
\(\beta=2^{24576}\) (quad-precision)
See Table 7.4.3.2, "Action for \(O E=1\)," on page 406 for trap-enabled Overflow exceptions.
See Table 7.4.4.2, "Action for UE=1," on page 411 for trap-enabled Underflow exceptions.
q The value defined in Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516, significand rounded to the target rounding precision, unbounded exponent range.
I The value defined in Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516, significand rounded to the target rounding precision, exponent bounded to the target rounding precision format exponent range,
error() The system error handler is invoked for the trap-enabled exception if MSR.FEO and MSR.FE1 are set to any mode other than the ignore-exception mode.

Table 59.VSX Scalar Floating-Point Final Result
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline Case & m

¢
Un
¢
d & U
O
Un
Un
d
d &  & \[
\begin{aligned}
& \mu \\
& \underset{N}{\tilde{U}} \\
& \dot{U} \\
& \frac{n}{u}
\end{aligned}
\] &  & \[
\begin{aligned}
& 0 \\
& \sigma \\
& \frac{\sigma}{4} \\
& 1 \\
& r_{0} \\
& \stackrel{1}{n} \\
& x \\
& >
\end{aligned}
\] & \[
\begin{aligned}
& 0 \\
& \sigma \\
& \frac{\sigma}{4} \\
& \tilde{c}^{\prime} \\
& \bar{x} \\
& \gg
\end{aligned}
\] & \[
\begin{gathered}
0 \\
\frac{0}{4} \\
\overline{1} \\
\bar{n} \\
\bar{x} \\
>
\end{gathered}
\] & \[
\begin{aligned}
& 0 \\
& \frac{0}{\sigma} \\
& \overline{4} \\
& \hline \bar{x} \\
& \bar{x} \\
& \gg
\end{aligned}
\] & \[
\begin{aligned}
& 0 \\
& 0 \\
& \frac{0}{4} \\
& 1 \\
& N^{1} \\
& N \\
& \underset{x}{x}
\end{aligned}
\] &  & \[
\begin{aligned}
& 0 \\
& 0 \\
& \overline{4} \\
& x^{\prime} \\
& N
\end{aligned}
\] &  & Is \(r\) incremented? \((|r|>|v|)\) &  &  & Returned Results and Status Setting \\
\hline \multirow{5}{*}{Normal} & - & - & - & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & no & - & - & - & T(r), class bfolr), filol , fr(0) \\
\hline & - & - & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & yes & no & - & - & \(\mathrm{T}(\mathrm{r})\), class_bfp(r), fi(1), fr(0), fx(x) \\
\hline & - & - & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & yes & yes & - & - & T(r), class_bfp(r), fi(1), fr(1), fx(x) \\
\hline & - & - & - & - & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & yes & no & - & - & T(r), class_bfo(r), fi(1), fr(0), fx(xx), errorl) \\
\hline & - & - & - & - & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & yes & yes & - & - & T(r), class_bfo(r), fi(l), fr(1), fx(xx), error() \\
\hline \multirow{5}{*}{Overflow} & - & 0 & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & - & - & - & - &  \\
\hline & - & 0 & - & - & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & - & - & - & - & T(r), class_bfp(r), fi(1), fr(?), fx(0x), fx(xX), error() \\
\hline & - & 1 & - & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & - & - & no & - & \(\mathrm{T}(\mathrm{q} \div \mathrm{\beta})\), class_bfp(q*B), fi(0), fr(0), fx(0x), error() \\
\hline & - & 1 & - & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & - & - & yes & no &  \\
\hline & - & 1 & - & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & - & - & yes & yes &  \\
\hline \multirow{8}{*}{Tiny} & - & - & 0 & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & no & - & - & - & T(r), class_bfolr), filol , fr (0) \\
\hline & - & - & 0 & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & yes & no & - & - & T(r), class_bfp(r), fi(l), fr(0), fx(ux), fx(xX) \\
\hline & - & - & 0 & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & yes & yes & - & - & \(\mathrm{T}(\mathrm{r})\), class_bfo(r), fi(1), fr(1), fx(ux), fx(xx) \\
\hline & - & - & 0 & - & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & yes & no & - & - & T(r), class_ ffp (r), fi(1), fr(0), fx(UX), fx(XX), error() \\
\hline & - & - & 0 & - & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & yes & yes & - & - & T(r), class_bfp(r), fi(1), fr(1), fx(UX), fx(XX), error() \\
\hline & - & - & 1 & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & - & - & no & - &  \\
\hline & - & - & 1 & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & - & - & yes & no &  \\
\hline & - & - & 1 & - & - & 0 & 0 & 0 & 0 & 0 & 0 & 0 & - & - & yes & yes &  \\
\hline \multicolumn{18}{|l|}{Explanation:} \\
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{}} & \multicolumn{16}{|l|}{The results do not depend on this condition.} \\
\hline & & \multicolumn{16}{|l|}{\begin{tabular}{l}
Places the result into the target VSR. \\
For scalar single-precision and double-precision results \\
VSR[XT]. dword[0] \& bfp_CONVERT_TO_BFP64(r) \\
VSR[XT].dword[1] \& OXUUUU_UUUU_UUUU_UUUU \\
For scalar quad-precision results \\
VSR[VRT+32] \&bf p_CONVERT_TO_BFP128(r)
\end{tabular}} \\
\hline \multicolumn{2}{|l|}{class_bfp(x)} & \multicolumn{16}{|l|}{\begin{tabular}{l}
Sets FPSCR. FPRF to the sign and class of \(x\). \\
FPSCR.FPRF \& fprf_CLASS_BFP32(x) (single-precision) \\
FPSCR.FPRF \(+\mathrm{fprf}^{-}\)CLASS \({ }^{-}\)BFP64(x) (double-precision) \\
FPSCR.FPRF \& fprf_CLASS_BFP128(x) (quad-precision)
\end{tabular}} \\
\hline \(f \mathrm{f}(\mathrm{x})\) & & \multicolumn{16}{|l|}{FPSCR. FX is set to 1 if \(F P S C R, X=0 . F P S C R, x\) is set to 1 .} \\
\hline \(f i(x)\) & & \multicolumn{16}{|l|}{FPSCR. FI is set to the value \(X\).} \\
\hline \(f r(x)\) & & \multicolumn{16}{|l|}{FPSCR. FR is set to the value \(x\).} \\
\hline \multirow[t]{2}{*}{\(\beta\)} & & \multicolumn{16}{|l|}{Wrap adjust
\[
\begin{array}{ll}
\beta=2^{192} & \text { (single-precision) } \\
\beta=2^{2536} & \text { (double-precision) } \\
\beta=2^{24576} & \text { (quad-precision) }
\end{array}
\]} \\
\hline & & \multicolumn{16}{|l|}{See Table 7.4.3.2, "Action for OE=1," on page 406 for trap-enabled Overflow exceptions. See Table 7.4.4.2, "Action for UE=1," on page 411 for trap-enabled Underflow exceptions.} \\
\hline q & & \multicolumn{16}{|l|}{The value defined in Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516, significand rounded to the target rounding precision, unbounded exponent range.} \\
\hline & & \multicolumn{16}{|l|}{The value defined in Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516, significand rounded to the target rounding precision, exponent bounded to the target rounding precision format exponent range.} \\
\hline error() & & \multicolumn{16}{|l|}{The system error handler is invoked for the trap-enabled exception if MSR.FEO and MSR.FE1 are set to any mode other than the ignore-exception mode.} \\
\hline
\end{tabular}

Table 59.VSX Scalar Floating-Point Final Result (Continued)

\section*{VSX Scalar Add Single-Precision XX3-form}
xsaddsp XT,XA,XB
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 60 & & T & & A & B & & 0 \\
\hline 0 & & \(6 \times B X X T X\) \\
\(2930 \mid 31\) \\
\hline
\end{tabular}
```

reset_xflags()
src1 \leftarrow VSR[32xAX+A].dword[0]
src2 \leftarrowVSR[32xBX+B].dword[0]
v }\leftarrow\mathrm{ AddDP(src1,src2)
result }\leftarrow\mathrm{ RoundToSP(RN,v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
vex_flag \leftarrow VE \& (vxsnan_flag | vxisi_flag)
if( ~vex_flag ) then do
VSR[32xTX+T].dword[0] \leftarrowConverSPtoDP(result)
VSR[32xTX+T].dword[1] \& 0xUUUU_UUUU_UUUU_UUUU
FPRF }\leftarrow\mathrm{ ClassSP(result)
FR \leftarrow inc_flag
FI }\leftarrowxx_fla
end
else do
FR}\leftarrow\textrm{ObO
FI }\leftarrow\textrm{ObO
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

Let src2 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X B]\).
src2 is added \({ }^{[1]}\) to src1, producing a sum having unbounded range and precision.

The sum is normalized \({ }^{[2]}\).
See Table 60, "Actions for xsaddsp," on page 520.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

\footnotetext{
1. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( \(G, R\), and \(X\) ) enter into the computation.
2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{}} & \multicolumn{8}{|c|}{src2} \\
\hline & & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \multirow{8}{*}{\[
\begin{aligned}
& \text { - } \\
& \text { נn }
\end{aligned}
\]} & -Infinity & \(v \leftarrow\)-nfinity & \(\mathrm{V} \leftarrow\) - Infinity & v ¢-Infinity & v ¢-Infinity & \(\mathrm{v} \leqslant-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{d} Q \mathrm{NaN}\) vxisi_flag \(\leftarrow 1\) & v ¢ SrC2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline & -NZF & \(v \leqslant-\) Infinity & \(\mathrm{V} \leftarrow A(\) src1, rrc 2\()\) & \(\mathrm{v} \leqslant \mathrm{SIC} 1\) & v ¢ STC1 & \(\mathrm{V} \leqslant \mathrm{A}(\mathrm{src} 1, \mathrm{src} 2)\) & \(v \leftarrow+\) Infinity & v ¢sic2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline & -Zero & \(v \leftarrow-\) Infinity & v ¢ STC2 & \(v \leftarrow\) Zero & \(v \leftarrow\) Rezd & v ¢ SrC2 & \(v \leftarrow+\) Infinity & \(v \leqslant\) src2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline & +Zero & \(v \leftarrow-\) Infinity & v ¢ SrC2 & \(v \leftarrow\) Rezd & \(v \leftarrow+\) Zero & \(\mathrm{v} \leqslant\) Src2 & \(v \leftarrow+\) Infinity & \(\mathrm{v} \leqslant\) Src2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline & +NZF & \(v \leqslant-\) Infinity & \(v \leftarrow A(\) src1,src2 \()\) & \(\mathrm{V} \leftarrow \mathrm{SrC} 1\) & v ¢ SrC1 & \(\mathrm{V} \leqslant \mathrm{A}\) (src1, src2) & \(v \leftarrow+\) Infinity & \(\mathrm{v} \leqslant\) src2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline & +Infinity & \(\mathrm{V} \leftarrow \mathrm{d}\) QNaN vxisi_flag \(\leftarrow 1\) & \(v \leftarrow+\) Infinity & \(v \leftarrow+\) Infinity & \(v \leftarrow+\) Infinity & \(v \leftarrow+\) Infinity & \(v \leftarrow+\) Infinity & v ¢ SrC2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline & QNaN & v ¢ S C 1 & v ¢ SC1 & v ¢ STC1 & v ¢ STC1 & v \& STC1 & v ¢ STC1 & v ¢ STC1 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \text { Src1 } \\
& \text { vxsnan_lag } \leftarrow 1
\end{aligned}
\] \\
\hline & SNaN & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
\mathrm{vxsnan} \text { flag } \leftarrow 1
\end{array}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1)\) vxsnan_flag \(\leftarrow 1\) & \(\mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1)\) vxsnan_flag \(\leftarrow 1\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
\mathrm{vxsnan} \text { flag } \leftarrow 1
\end{array}
\] & \(\mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1)\) vxsnan_flag \(\leftarrow 1\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{Explanation:} \\
\hline srcl & The double-precision floating-point value in doubleword element 0 of VSR[ XA] . \\
\hline sicl & The double-precision floating-point value in doubleword element 0 of VSR [ XB]. \\
\hline dQNaN & Default quiet \(\mathrm{NaN}(0 \times 7 \mathrm{FF} 8\) _0000_0000_0000). \\
\hline NZF & Nonzero finite number. \\
\hline Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). \\
\hline A \((x, y)\) & Return the normalized sum of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. Note: If \(x=\cdot y, v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline \(Q(x)\) & Return a QNaN with the payload of \(x\). \\
\hline \(v\) & The intermediate result having unbounded signficand precision and unbounded exponent range. \\
\hline
\end{tabular}

Table 60.Actions for xsaddsp

VSX Scalar Add Quad-Precision [using round to Odd] X-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline xsaddq & & T,VR & VRB & & & \\
\hline xsadd & & RT,VR & ,VRB & & & \\
\hline \[
06
\] & \[
{ }_{6} \text { VRT }
\] & \[
{ }_{11} \text { VRA }
\] & \[
{ }_{16} \text { VRB }
\] & 21 & 4 & RO2 \\
\hline
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable(l
reset_xflags()
srCl \&bfp_CONVERT_FROM_BFP128(VSR[VRA+32])
srcl \leftarrowbfp_CONVERT_FROM_BFP128(VSR[VRB+32])
v \leftarrowbfp_ADD(srci, srici)
rnd }\leftarrow\mathrm{ bfp_ROUND_TO_BFP128(RO, FPSCR,RN, v)
result \& bfp_CONVERT_TO_BFP128(rnd)
if(vxsnan_flag) then SetFX(FPSCR.VXSNAN)
if(vxisi_flag) then SetFX(FPSCR.VXISI)
if(Ox_flag) then SetFX(FPSCR.OX)
if(ux_flag) then SetFX(FPSCR.UX)
if(xx_flag) then SetFX(FPSCR.XX)
vxflag \leftarrowvxsnan_flag| vxisi_flag
ex_flag \leftarrowFPSCR.VE\&vx_flag
if ex_flag=0 then do
VSR[VRT+32] \leftarrowresult
FPSCR.FPRF}\leftarrowf\mathrm{ frf_CLASS_BFP128(result)
end
FPSCR,FR\leftarrow(vx_flag=0) \& inc_flag
FPSCR,Fl}\leftarrow(v\mp@subsup{x}{-}{-}f|ag=0)\& XX_ flag

```

Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let \(\operatorname{src} 2\) be the floating-point value in VSR[VRB +32 ] represented in quad-precision format.

If either srcl or srcl is a Signalling NaN, an Invalid Operation exception occurs and VXSNAN is set to 1.

If srCl and \(\mathrm{srC2}\) are Infinity values having opposite signs, an Invalid Operation exception occurs and VXI SI is set to 1 .

If srcl is a Signalling NaN , the result is the Quiet NaN corresponding to srcl.

Otherwise, if srcl is a Quiet NaN , the result is srcl .
Otherwise, if \(\mathrm{src2}\) is a Signalling NaN , the result is the Quiet NaN corresponding to src2.

Otherwise, if src 2 is a Quiet NaN , the result is src .
Otherwise, if srcl and \(\mathrm{srcl}_{2}\) are Infinity values having opposite signs, the result is the default Quiet \(\mathrm{NaN}^{[1]}\).

Otherwise, do the following.
The normalized sum of srcl added to srcl is produced with unbounded significand precision and exponent range.

See Table 61, "Actions for xsaddqp[0]," on page 522.

If the intermediate result is Tiny (i.e., the unbiased exponent is less than -16382) and \(U E=0\), the significand is shifted right \(N\) bits, where \(N\) is the difference between - 16382 and the unbiased exponent of the intermediate result. The exponent of the intermediate result is set to the value - 16382.

If \(R O=1\), let the rounding mode be Round to Odd. Otherwise, let the rounding mode be specified by RN. The intermediate result is rounded to quad-precision using the specified rounding mode.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into VSR[ VRT +32 ] in quad-precision format.

FPRF is set to the class and sign of the result. FR is set to indicate if the rounded result was incremented. Fl is set to indicate the result is inexact.

If a trap-disabled Invalid Operation exception occurs, \(F R\) and \(F I\) are set to 0 .

If a trap-enabled Invalid Operation exception occurs, VSR[VRT+32] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered:}

FPRF FR FI
FX VXSNAN VXISI OX UX XX

\section*{VSR Data Layout for xsaddqp[o]}

VSR[ VRA+32]
\(\square\)
VSR[VRB+32]


\footnotetext{
1. The quad-precision default Quiet NaN is the value, \(0 \times 7 \mathrm{FFF} \mathrm{C}_{8} 8000_{-} 0000_{-} 0000_{-} 0000_{-} 0000 \_0000\).
}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{}} & \multicolumn{8}{|c|}{src2} \\
\hline & & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \multirow{8}{*}{\[
\underset{\sim}{0}
\]} & -Infinity & \multirow[t]{5}{*}{\(v \leftarrow\) Infinity} & & & & & \[
\begin{gathered}
v \leftarrow \operatorname{dQNaN} \\
\text { vxisi_flag } \leftarrow 1 \\
\hline
\end{gathered}
\] & & \\
\hline & -NZF & & \(v \leftarrow \operatorname{adod}(\operatorname{sicl}, \operatorname{sic} 2)\) & & & \(v \leftarrow \operatorname{add}(\operatorname{scc}\), scicl \()\) & & & \\
\hline & -Zero & & & \(V \leftarrow\) Zero & \(v \leqslant\) Rezd & & & & \\
\hline & +Zero & & & \(v \leftarrow\) Rezd & \(v \leftarrow+\) Eero & & & & \(v \leftarrow\) quiet \(\left(s\right.\) sch \(\left.^{2}\right)\) vxsnan_flag \(\leftarrow 1\) \\
\hline & +NZF & & \(v \leftarrow \operatorname{add}(\operatorname{sicl}, \operatorname{sic} 2)\) & V & & \(v \leftarrow \operatorname{add}(\operatorname{scc}\), scicl \()\) & & & \\
\hline & +Infinity & \[
\begin{gathered}
v \leftarrow \text { dQNaN } \\
\text { vxisi flag }
\end{gathered}
\] & & & & & \(v \leftarrow\) Infinity & & \\
\hline & QNaN & \multicolumn{5}{|c|}{\(v \leftarrow\) sicl} & & & \[
\begin{gathered}
v \leftarrow \operatorname{sic} 1 \\
\text { vxsnan_flag } \leftarrow 1
\end{gathered}
\] \\
\hline & SNaN & \multicolumn{8}{|c|}{\[
\begin{aligned}
& v \leftarrow \text { quiet (sicl) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\]} \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{Explanation:} \\
\hline sicl & The quad-precision floating-point value in VSR [ VRA +32\(]\). \\
\hline sic2 & The quad-precision floating-point value in VSR [ VRB +32\(]\). \\
\hline dQ NaN &  \\
\hline NZF & Nonzero finite number. \\
\hline Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude and opposite signs). \\
\hline add ( \(x, y\) ) & The floating-point value \(y\) is added \({ }^{1}\) to the floating-point value \(x\). Return the normalized \({ }^{2}\) sum, having unbounded significand precision and exponent range. \\
\hline & When \(x=\cdot y, v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline quiet ( x ) & Convert x to the corresponding Quiet NaN by setting the most significant fraction bit to 1. \\
\hline v & The intermediate result having unbounded significand precision and unbounded exponent range. \\
\hline
\end{tabular}

\section*{| Table 61. Actions for xsaddqp[o]}
1. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate difference.
2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.

\section*{VSX Scalar Compare Exponents Double-Precision XX3-form}
xscmpexpdp BF,XA,XB

```

if MSR.VSX=O then VSX_Unavailable()
srcl \&VSR[32xAX+A].dword[0]
scc2 \&vSR[32xBX+B],dword[0]
src1.exponent \& ExTz(src1., bit[1:11])
src2.exponent \leftarrow ExTZ(scr(2.bit[1:11])
scl.fraction < srcl.bit[12:63]
src2.fraction\leftarrow src2.bit[12:63]
srcl.class.NaN \leftarrow(srcl.exponent = 2047) \& (srcl.fraction!= 0)
src2.class.NaN }\leftarrow(scc2.exponent =2047) \& (src2.fraction!=0
It_flag}\leftarrow(\mathrm{ srcl.exponent < src2, exponent)
gt_flag}\leftarrow(\mathrm{ srcl, exponent > src2, exponent )
eq_ flag}\leftarrow(\mathrm{ srcl.exponent = src2, exponent
vo_flag}\leftarrow\mathrm{ srcl.class.NaN | sic2.class.NaN
CR. bit[4\timesBF+32]}\leftarrow FPSCR.FL\leftarrow!vo_flag\& |t_flag
CR.bit[4\timesBF+33]}\leftarrowFPSCR.FG\leftarrow!uo_flag \& gt_flag
CR. bit[4xBF+34]}\leftarrowFPSCR.FE \leftarrow!uo_flag\& eq_fla

```

| VSR Data Layout for xscmpexpdp
I
\begin{tabular}{|c|c|c|c|}
\hline srcl & VSR[ XA]. dword[0] & unused & \\
\hline \multirow[t]{2}{*}{srcz} & VSR[ XB]. dword[0] & unused & \\
\hline & 0 & 64 & 127 \\
\hline
\end{tabular}

\section*{Version 3.0}

\section*{VSX Scalar Compare Exponents Quad-Precision X-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline xscmpe & & & F,VRA & RB & & & \\
\hline 063 & \({ }_{6} \mathrm{BF}\) & // & 11 VRA & \({ }_{16}\) VRB & 21 & 164 & | 1 \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailable()
reset_flags()
sicl \(\leftarrow\) VSR[VRA +32\(]\)
\(\operatorname{srC2} \leftarrow V S R[V R B+32]\)
srcl.exponent \(\leftarrow\) EXTZ(srcl.bit[1:15])
src2. exponent \(\leftarrow\) EXTZ(src2.bit [1:15])
srcl.fraction \(\leftarrow\) EXTZ(srcl.bit[16:127])
src2.fraction \(\leftarrow\) EXTZ(src2.bit [16:127])
srcl.class. \(\mathrm{NaN} \leftarrow(\) srcl.exponent \(=32767) \&(\) srcl.fraction \(!=0)\)
src2.class. \(\mathrm{NaN} \leftarrow(\) src2.exponent \(=32767) \&(\) src2.fraction \(!=0)\)
It_flag \(\leftarrow\) (srcl.exponent < src2. exponent \()\)
gt_flag \(\leftarrow\) (srcl.exponent \(>\) srcl. exponent \()\)
eq_flag \(\leftarrow(\) srcl.exponent \(=\) src2. exponent \()\)
uo_flag \(\leftarrow\) srcl.class.NaN \(\mid\) srcl.class.NaN
CR. bit \([4 \times B F+32] \leftarrow\) FPSCR.FL \(\leftarrow!\) Uo_flag \& It filag
CR. bit \([4 \times B F+33] \leftarrow F P S C R . F G \leftarrow!u 0_{-} f l a g \& g t\) flag
CR. bit \([4 \times B F+34] \leftarrow F P S C R . F E \leftarrow!\) uo_flag \& eq_flag
CR. bit \([4 \times B F+35] \leftarrow\) FPSCR.FU \(\leftarrow U 0^{-} f 1\) ag

Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let src2 be the floating-point value in VSR[VRB+32] represented in quad-precision format.

The exponent of \(\operatorname{srcl}\) is compared with the exponent of src2 as unsigned integer values. The result of the compare is placed into \(F P C C\) and \(C R\) field \(B F\).

Special Registers Altered:
CR field BF
FPCC
| VSR Data Layout for xscmpexpqp
I VSR[VRA+32]
I \(\square\)
I VSR[VRB +32\(]\)
I
sic2

VSX Scalar Compare Equal Double-Precision XX3-form

\section*{xscmpeqdp XT,XA,XB}

```

if MSR.VSX=O then VSX_Unavailablell
srcl \leftarrow bf p_CONVERT_FROM_BFP64(VSR[ 32xAX+A].dword[0])
src2 \leftarrow bfp_CONVERT_FROM_BFP64(VSR[32xBX+B].dword[0])
vxsnan_flag \leftarrow(srcl.class="SNaN") | (src2.class="SNaN")
vex_flag \leftarrow FPSCR.VE \& vxsnan_flag
if(vxsnan_flag) SetFX(FPSCR.VXSNAN)
if (vex_flag=0) then do
if bf _COMPARE_EQ(srcl, src2)=1 then
VSR[32XTX+T].dword[0]}\leftarrow0XFFFFFFFFFFFFFFFF
VSR[32xTX+T].dword[1]}\leftarrow0\times0000_0000_0000_000
end
else do
VSR[32xTX+T].dword[0]}\leftarrow0\times0000_0000_0000_000
VSR[32xTX+T].dword[1]}\leftarrow0x0000_0000_0000_0000
end
end

```

Let \(X T\) be the value \(32 \times T X+T\).

Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(\mathrm{srcl}_{1}\) be the double-precision floating-point value in doubleword 0 of VSR[ XA].

Let src2 be the double-precision floating-point value in doubleword 0 of VSR[ XB].

If srcl or srcl is a SNaN , an Invalid Operation exception occurs.
srcl is compared to srcl .
A NaN compared to any value, including itself, compares false for the predicate, equal.

The contents of doubleword 0 of VSR[XT] are set to \(O_{X F F F F}\) FFFF_FFFF_FFFF if \(\operatorname{srcl}\) is equal to src2, and are set to \(0^{-} \times 0000_{-}^{-} 0000^{-} 0000 \_0000\) otherwise.

The contents of doubleword 1 of VSR[ XT] are set to \(0 \times 0000 \_0000 \_0000 \_0000\).

If a trap-enabled Invalid Operation occurs, VSR[ XT] is not modified.

\section*{Special Registers Altered:}

FX VXSNAN

\section*{Version 3.0}

\section*{VSX Scalar Compare Greater Than or Equal Double-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline xscmpg & \multicolumn{4}{|c|}{XT, XA, XB} & \\
\hline \[
00
\] & \[
6 \quad \text { T }
\] & \[
11 \quad \mathrm{~A}
\] & \[
{ }_{16} \text { B }
\] & \[
\begin{array}{ll} 
& 19 \\
21 &
\end{array}
\] &  \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailable(l)
sicl \(\leftarrow\) bf p_CONVERT_FROM_BFP64(VSR[ \(32 \times A X+A]\). dword[ 0\(])\)
sicl \(\leftarrow\) bf P_CONVERT_FROM_BFP64(VSR[ \(32 \times B X+B]\).dword[0])
if (srcl.class="SNaN") | (src2.class="SNaN") then do vxsnan_flag \(\leftarrow 0 b 1\) if(FPSCR.VE=O) then vxVC_flag \(\leftarrow 0 b 1\)
end
else
vxvc_flag \(\leftarrow(\) srcl.class \(=\) " ONaN" \() \mid\) (src2.class="ONaN")
vex_flag \(\leftarrow\) FPSCR.VE \& (vxsnan_flag \(\mid\) vxvc_flag)
if (vxsnan_flag=1) SetFX(FPSCR.VXSNAN)
if (VXCV_flag=1) SetFX(FPSCR.VXVC)
if (vex_flag=0) then do
if \(b \bar{f} p\) _COMPARE_GE(srcl \(s(c 2)=1\) then

VSR[32xTX+T].dword[1] \(\leftarrow 0 \times 0000 \_0000 \_0000 \_0000\)
end
else do
VSR[ \(32 \times T X+T]\).dword \([0] \leftarrow 0 \times 0000 \_00000000 \_0000\)
\(\operatorname{VSR}[32 \times T X+T]\), dword [1] \(\leftarrow 0 \times 0000^{-} 0000^{-} 0000^{-} 0000\)
end
end
I

\section*{VSX Scalar Compare Greater Than Double-Precision XX3-form}
xscmpgtdp \(\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)

```

if MSR.VSX=O then VSX_Unavailable()
srcl \leftarrow bfp_CONVERT_FROM_BFP64(VSR[ 32xAX+A].dword[0])
src2 \leftarrow bfp_CONVERT_FROM_BFP64(VSR[32\timesBX+B].dword[0])
if (srcl.class="SNaN")| (src2.class="SNaN") then do
vxsnan_flag \leftarrow Obl
if(FPSCR.VE=0) then vxvc_flag \leftarrow Obl
end
else
vxvc_flag \leftarrow (srcl.class="ONaN")| (src2.class="QNaN")
vex_flag \leftarrow FPSCR.VE \& (vxsnan_flag| vxvc_flag)
if (vxsnan_flag=1) SetFX(FPSCR. VXSNAN)
if (vxcv_flag=1) SetFX(FPSCR.VXVC)
if (vex_flag=0) then do
if bfp_COMPARE_GT(srcl, src2)=1 then
VSR[32xTX+T].dword[0] \leftarrowOXFFFF_FFFF_FFFF_FFFF
VSR[32xTX+T].dword[1] \leftarrow0x0000_0000_0000_0000
end
else do
VSR[32xTX+T].dword[0]}\leftarrow0\times0000_0000_0000_000
VSR[32xTX+T].dword[1]}\leftarrow0\times0000-0000_0000-000
end
end

```
I

\section*{Version 3.0}

\section*{VSX Scalar Compare Not Equal \\ Double-Precision XX3-form}
Xscmpnedp \(\mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)
\begin{tabular}{|c|cc|c|c|c|c|c|c|}
\hline 60 & & T & A & B \\
\hline 0 & & 6 & & 11 & & 16 & & 21 \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailable(l)
srcl \(\leftarrow\) bf p_CONVERT_FROM_BFP64(VSR[ \(32 x A X+A]\). dword[ 0\(])\)
src2 \(\leftarrow\) bfp_CONVERT_FROM_BFP64(VSR[ \(32 \times B X+B]\). dword[0])
vxsnan_flag \(\leftarrow(\) srcl.class =" SNaN") \(\mid\) (srcl.class="SNaN")
if(vxsnan_flag) SetFX(FPSCR.VXSNAN)
vex_flag \(\leftarrow\) FPSCR.VE \& vxsnan_flag
if (vex_fag=0) then do
if bfo_COMPARE NE (srcl, src2) \(=1\) then
\(\operatorname{VSR}[32 \times T X+T]\). \(d\) word \([0] \leftarrow\) OXFFFF_FFFF_FFFF_FFFF \(^{2}\) VSR[32xTX+T].dword[1] \(\leftarrow 0 \times 0000\) _0000_0000_0000
end
else do
VSR[ \(32 \times T X+T]\).dword \([0] \leftarrow 0 \times 0000 \_0000 \_0000 \_0000\)
VSR \([32 \times T X+T]\), dword [1] \(\leftarrow 0 \times 0000^{-} 0000^{-} 0000^{-} 0000\)
end
end
Let XT be the value \(32 \times T X+T\).

Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(\operatorname{srcl}\) be the double-precision floating-point value in doubleword 0 of VSR[ XA].

Let \(\mathrm{src}_{2}\) be the double-precision floating-point value in doubleword 0 of VSR[ XB].
srcl is compared to src 2 .
A NaN compared to any value, including itself, compares true for the predicate, not equal.

The contents of doubleword 0 of VSR[XT] are set to \(O_{X F F F F}\) FFFF FFFF_FFFF if SrCl is not equal to \(\mathrm{SrC2}\), and are set to \(0 \times 0000-0000 \_0000 \_0000\) otherwise.

The contents of doubleword 1 of VSR[XT] are set to \(0 \times 0000 \_0000 \_0000-0000\).

If a trap-enabled Invalid Operation occurs, VSR[XT] is not modified.

Special Registers Altered:
FX VXSNAN

\section*{VSX Scalar Compare Ordered Double-Precision XX3-form}


Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(\operatorname{srcl}\) be the double-precision floating-point value in doubleword element 0 of VSR[ XA].

Let \(\operatorname{src}\) c be the double-precision floating-point value in doubleword element 0 of VSR[ XB].
srcl is compared to srcl .
Zeros of same or opposite signs compare equal.
Infinities of same signs compare equal.
See Table 62, "Actions for xscmpodp - Part 1: Compare Ordered," on page 530.

The result of the compare is placed into CR field \(B F\) and the FPCC.

If either of the operands is a NaN , either quiet or signaling, CR field BF and the FPCC are set to reflect unordered. If either of the operands is a Signaling NaN, VXSNAN is set, and Invalid Operation is disabled \((V E=0), V X V C\) is set. If neither operand is a Signaling NaN but at least one operand is a Quiet NaN, VXVC is set.

See Table 63, "Actions for xscmpodp - Part 2: Result," on page 530.

\section*{Special Registers Altered}

CR field BF
FPCC FX VXSNAN VXVC
VSR Data Layout for xscmpodp
srcl \(=\) VSR[ XA]
\begin{tabular}{|r|l|}
\hline \multicolumn{2}{|c|}{ DP } \\
\hline SrC2 \(=\) VSR[ XB] \\
\hline DP & unused \\
\hline 0 & 64 \\
\hline
\end{tabular}

\section*{Programming Note}

This instruction can be used to operate on single-precision source operands.
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & cc \(<060010\) & cc<0b1000 & cc¢ \(<0 b 1000\) & cc¢0b1000 & cc<0b1000 & cc \(<0 b 1000\) & cc \(\leftarrow 0 b 0001\) vxvc_flag\& 1 & cc \(-0 b 0001\) vxsnan_flag -1 vxvc__lag\&(VE=0) \\
\hline -NZF & cc¢ \(<060100\) & ccte (srct , src2) & cc¢ \(<0 b 1000\) & \(c c \leftarrow 0 b 1000\) & \(c c \leftarrow 0 b 1000\) & cc \(<0 b 1000\) & cc \(\leftarrow 0 b 0001\) vxvc_flag \(\leftarrow 1\) & cc \(\leqslant 0 b 0001\) vxsnan_flag -1 vxvc_lage(VE=0) \\
\hline -Zero & cc¢ \(-0 b 0100\) & cc¢ \(<060100\) & cc \(<060010\) & \(c \mathrm{c} \leftarrow 0 \mathrm{~b} 0010\) & cc \(<\) Ob1000 & cc \(<0 b 1000\) & cc \(<0 b 0001\) vxvc_flags-1 & cc \(-0 b 0001\) vxsnan_flag -1 vxvc__lagヶ(VE=0) \\
\hline - +Zero & cc \(<060100\) & \(c \mathrm{c} \leftarrow 0 \mathrm{Ob} 100\) & cc<0b0010 & cc \(<060010\) & cc \(<0 b 1000\) & cc \(<0 b 1000\) & cc \(\leftarrow 0 b 0001\) vxvc_flag \(\leftarrow 1\) & cc \(-0 b 0001\) vxsnan_flag -1 vxvc flage(VE=0) \\
\hline +NZF & cç-0b0100 & cc¢ \(<060100\) & cc \(<060100\) & cc< \(<060100\) & \(\mathrm{cc} \leftarrow \mathrm{C}(\) srcc 1 ,src2) & \(c c<0 b 1000\) & \[
c c \leftarrow 0 b 0001
\]
|vxvc_flag<1 & cc \(<0 b 0001\) vxsnan_flag -1 vxvc_flag↔(VE=0) \\
\hline +Infinity & cc \(<060100\) & \(c \mathrm{c} \leftarrow 0 \mathrm{Ob} 100\) & cc \(<060100\) & \(c \mathrm{c} \leftarrow 0 \mathrm{Ob} 100\) & cc \(<060100\) & cc \(<060010\) & cc \(-0 b 0001\) vxvc_lag_1 & cc \(-0 b 0001\) vxsnan_flag \(<1\) vxvc__lag↔(VE=0) \\
\hline QNaN & cc \(\leftarrow 0 b 0001\) vxvc_flag<1 & cc \(<0 b 0001\) vxvc_flag<1 & cc \(\leftarrow 0 b 0001\) vxvc_flag\& & cc \(<0 b 0001\) vxvc_flag\&1 & cc \(<0 b 0001\) vxvc_flage1 & cc \(\leftarrow 0 b 0001\) vxvc_flag<1 & cc \(\leftarrow 0 b 0001\) vxvc_flag \(\leftarrow 1\) & cc \(<0 b 0001\) vxsnan_flag -1 vxvc flage-(VE=0) \\
\hline SNaN & cc \(-0 b 0001\) vxsnan_flag -1 vxvc_flag↔(VE=0) & cc \(-0 b 0001\) vxsnan_flag\& 1 vxv__lag\&(VE=0) & cc \(-0 b 0001\) vxsnan_flag -1 vxvc_flag\&(VE=0) & cc \(<0 b 0001\) vxsnan_flag\&-1 vxvc_flag\&(VE=0) & cc \(\leftarrow 0 b 0001\) vxsnan_flag -1 vxv__flag↔(VE=0) & cc \(-0 b 0001\) vxsnan_flag\& 1 vxvc_flag\&(VE=0) & cc \(-0 b 0001\) vxsnan_flag<1 vxv__flag\&(VE=0) & cc \(\leftarrow 0 b 0001\) vxsnan_flag\& 1 vxv__lag\&(VE=0) \\
\hline
\end{tabular}

\section*{Explanation:}
\begin{tabular}{|c|c|}
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
\hline src2 & The double-precision floating-point value in doubleword element 0 of VSR[XB]. \\
\hline NZF & Nonzero finite number. \\
\hline C(x,y) & \begin{tabular}{l}
The floating-point value \(x\) is compared to the floating-point value \(y\), returning one of three 4 -bit results. \\
0b1000 when x is greater than y \\
0b0100 when \(x\) is less than \(y\) \\
Ob0010 when \(x\) is equal to \(y\)
\end{tabular} \\
\hline cc & The 4-bit result compare code. \\
\hline
\end{tabular}

\section*{Table 62.Actions for xscmpodp - Part 1: Compare Ordered}


Table 63.Actions for xscmpodp - Part 2: Result

VSX Scalar Compare Ordered Quad-Precision X-form
\[
\begin{aligned}
& \text { xscmpoqp } \\
& \begin{array}{|l|l|l|l|l|l|l|}
\hline 63 & \begin{array}{l}
\text { BF }
\end{array} \\
\hline 0 & \begin{array}{l}
\text { I/ } \\
9
\end{array} & \text { VRA } \\
\hline
\end{array} \\
& \hline
\end{aligned}
\]
```

if MSR.VSX=O then VSX_Unavailable(l
reset_xflags()
srcl \&bfp CONVERT FROM BFP128(VSR[VRA+32])
src2 \&bfp_CONVERT_FROM_BFP128(VSR[VRB+32])
if( srcl.class.SNaN | srcl.class.SNaN ) then do
vxsnan_flag < Obl
if(FPSCR,VE=O) then vxvc_flag \& Obl
end
else if( srcl.class, QNaN | srcl.class.QNaN ) then vxvc_flag \leftarrow Obl
cc.bit[0] \& bfp_COMPARE_LT(srcl,src2)
cc,bit[1]}\leftarrowbfp COMPARE GT(srcl,src2
cc.bit[2]}\leftarrow\mathrm{ bf_ COMPARE_EQ(srcl,src2)
cc.bit[3] \leftarrow srcl.class.SNaN |rcl.class.QNaN
srcl.class.SNaN| srcl.class, QNaN

```
if(vxsnan flag) then Set FX(FPSCR. VXSNAN)
if(vxvc_fiag) then SetFX(FPSCR.VXVC)
FPSCR.FPCC \(\leftarrow C\)
CR.field \([B F] \leftarrow C C\)

Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let src2 be the floating-point value in VSR[VRB+32] represented in quad-precision format.
srcl is compared to srcl .

Zeros of same or opposite signs compare equal Infinities of same signs compare equal.

Bit 0 of \(C R\) field \(B F\) and \(F L\) are set to indicate if \(\mathrm{srCl}^{\text {is }}\) less thansic2.

Bit 1 of \(C R\) field \(B F\) and \(F G\) are set to indicate if \(s r C 1\) is greater than sic2.

Bit 2 of \(C R\) field \(B F\) and \(F E\) are set to indicate if \(S r C 1\) is equal to sicl.

Bit 3 of \(C R\) field \(B F\) and \(F U\) are set to indicate unordered (i.e., srcl orsrcl is a NaN ).

If either of the operands is a NaN , either quiet or signaling, CR field BF and the FPCC are set to reflect unordered. If either of the operands is a Signaling NaN , an Invalid Operation exception occurs and VXSNAN is set, and if Invalid Operation exceptions are disabled ( \(V E=0\) ), VXVC is set. If neither operand is a Signaling NaN but at least one operand is a Quiet NaN, an Invalid Operation exception occurs and VXVC is set.

\section*{Special Registers Altered:}

CR field BF
FPCC FX VXSNAN VXVC

VSR Data Layout for xscmpoqp
VSR [ VRA+32]


VSR[ VRB+32]

\section*{Version 3.0}

\section*{VSX Scalar Compare Unordered} Double-Precision XX3-form
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{xscmpudp BF,XA, XB} \\
\hline 00 & \({ }_{6} \mathrm{BF}\) & & 11. & 16 & 21 & 35 &  \\
\hline \multicolumn{8}{|l|}{\(X A \quad \leftarrow A X \| A\)} \\
\hline \multicolumn{8}{|l|}{\(X B \leqslant B X \| B\)} \\
\hline \multicolumn{8}{|l|}{reset_xflags ()} \\
\hline \multicolumn{8}{|l|}{srcl \(\leftarrow \operatorname{VSR}[X A]\{0: 63\}\)} \\
\hline \multicolumn{8}{|l|}{SrCL \(\leftarrow \operatorname{VSR}[X B]\{0: 63\}\)} \\
\hline \multicolumn{8}{|l|}{if( IssNaN(srcl) | IssNaN(src2) ) then vxsnan_flag \(\leftarrow 1\)} \\
\hline \multicolumn{8}{|l|}{\(\mathrm{FL} \quad \leftarrow\) Compareltdp(srcl, src2)} \\
\hline \multicolumn{8}{|l|}{FG \(\leftarrow\) CompareGTDP(srcl, src2)} \\
\hline \multicolumn{8}{|l|}{FE \(\leftarrow\) CompareEQDP(srcl,src2)} \\
\hline \multicolumn{8}{|l|}{FU ( \(\leftarrow\) |sNAN(srcl) | | sNAN(src2)} \\
\hline \multicolumn{8}{|l|}{\(C R[B F] \leftarrow F L\|F G\| F E \| F U\)} \\
\hline \multicolumn{8}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline
\end{tabular}

Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src1 is compared to src2.
Zeros of same or opposite signs compare equal equal.
Infinities of same signs compare equal.
See Table 64, "Actions for xscmpudp - Part 1: Compare Unordered," on page 533.

The result of the compare is placed into CR field BF and the FPCC.

If either of the operands is a NaN , either quiet or signaling, CR field BF and the FPCC are set to reflect unordered. If either of the operands is a Signaling \(\mathrm{NaN}, \mathrm{VXSNAN}\) is set.

See Table 65, "Actions for xscmpudp - Part 2: Result," on page 533.

\section*{Special Registers Altered}

CR[BF]
FPCC FX VXSNAN

\section*{Programming Note}

This instruction can be used to operate on single-precision source operands.


Table 64.Actions for xscmpudp - Part 1: Compare Unordered
\begin{tabular}{|c|c|c|c|}
\hline \[
\boldsymbol{>}
\] & \[
\begin{gathered}
\frac{0}{\pi} \\
\frac{\pi}{4} \\
\frac{1}{\pi} \\
\frac{1}{\omega} \\
\times \\
>
\end{gathered}
\] & & rned Results a \\
\hline - & 0 & FPC & \(\leftarrow \mathrm{cc}, \mathrm{CR}[\mathrm{BF}] \leftarrow \mathrm{cc}\) \\
\hline 0 & 1 & FPC & \(\leftarrow c \mathrm{c}, \mathrm{CR}[\mathrm{BF}] \leftarrow \mathrm{cc}, \mathrm{fx}\) \\
\hline 1 & 1 & & \(\leftarrow c c, C R[B F] \leftarrow c c, f x\) \\
\hline \multicolumn{3}{|l|}{Explanation} & \\
\hline \multicolumn{2}{|l|}{} & & The results do no \\
\hline \multicolumn{2}{|r|}{cc} & & The 4-bit result a \\
\hline \multicolumn{2}{|r|}{fx(x)} & & FX is set to 1 if x \\
\hline \multicolumn{2}{|r|}{error()} & & The system erro to any mode oth \\
\hline \multicolumn{2}{|r|}{FX} & & Floating-Point Su \\
\hline \multicolumn{3}{|c|}{VXSNAN} & Floating-Point Inv \\
\hline
\end{tabular}

Table 65.Actions for xscmpudp - Part 2: Result

\section*{Version 3.0}

\section*{VSX Scalar Compare Unordered Quad-Precision X-form}
xscmpuqp
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 63 & BF & \begin{tabular}{l} 
I/ VRA, VRB \\
9
\end{tabular} & VRA & VRB & & 644 \\
\hline 0 & & 6 & & 16 & 21 & \\
31 \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailable(l)
reset_xflags()
STCl \(\leftarrow\) bfo_CONVERT_FROM_BFP128(VSR[VRA+32])
\(\operatorname{srC2} \leftarrow b \mathrm{P}_{-}^{-}\)CONVERT_FROM_BPP128(VSR[VRB+32]) \(^{-}\)
vxsnan_flag \(\leftarrow\) srcl.class. SNaN \(\mid\) srcl.class. SNaN
cc.bit[0] \(\leftarrow\) bfo_COMPARE_LT(srcl,src2)
cc.bit[1] \(\leftarrow\) bfp_COMPAREGT(srcl,srcz)
cc.bit[2] \(\leftarrow\) bf p_COMPARE_EQ(sicl,sic2)
cc.bit[3] \(\leftarrow\) srci.class. \(\overline{\text { SNaN }}\) | sicl.class. QNaN | srcl.class.SNaN | srcl.class.QNaN
if(vxsnan_flag) then SetFX(FPSCR. VXSNAN)
FPSCR.FPCC \(\leftarrow C C\)
\(C R\), fiel \(d[B F] \leftarrow C C\)

Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let \(\operatorname{srcl}\) be the floating-point value in VSR[VRB +32 ] represented in quad-precision format.
srcl is compared to src 2 .
Zeros of same or opposite signs compare equal. Infinities of same signs compare equal.

Bit 0 of CR field BF and FL are set to indicate if SrCl is less thansic2.

Bit 1 of \(C R\) field \(B F\) and \(F G\) are set to indicate if \(s r c 1\) is greater than \(\mathrm{SrC2}\).

Bit 2 of \(C R\) field \(B F\) and \(F E\) are set to indicate if SrCl is equal to src2.

Bit 3 of \(C R\) field \(B F\) and \(F U\) are set to indicate unordered (i.e., srcl or srcl is a NaN ).

If either of the operands is a Signaling NaN, an Invalid Operation exception occurs and VXSNAN is set to 1.

\section*{Special Registers Altered:}

CR field BF
FPCC FX VXSNAN
VSR Data Layout for xscmpuqp
| VSR[VRA+32]
I


VSR[ VRB+32]
I

VSX Scalar Copy Sign Double-Precision XX3-form

\begin{tabular}{ll}
XT & \(\leftarrow \mathrm{TX} \| \mathrm{T}\) \\
XA & \(\leftarrow \mathrm{AX} \| \mathrm{A}\) \\
XB & \(\leftarrow \mathrm{BX} \| \mathrm{B}\) \\
result \(\{0: 63\}\) & \(\leftarrow \operatorname{VSR}[\mathrm{XA}]\{0\} \|\) VSR [XB] \{1:63\} \\
\(\operatorname{VSR}[\mathrm{XT}]\) & \(\leftarrow\) result \(\|\) 0xUUUU_UUUU_UUUU_UUUU
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Bit 0 of \(\operatorname{VSR}[X T]\) is set to the contents of bit 0 of VSR[XA].

Bits 1:63 of VSR[XT] are set to the contents of bits 1:63 of VSR[XB].

The contents of doubleword element 1 of VSR[XT] are undefined.
```

Special Registers Altered
None
VSR Data Layout for xscpsgndp
srcl = VSR[ XA]

| DP | unused |
| :---: | :---: |

srC2 = VSR[XB]

| DP | unused |
| :---: | :---: |

tgt = VSR[XT]

| DP | undefined |
| :--- | :--- |
| 0 | 64 |

```

\section*{Programming Note}

This instruction can be used to operate on single-precision source operands.

VSX Scalar Copy Sign Quad-Precision X-form
xscpsgnqp VRT,VRA,VRB

if MSR.VSX=O then VSX_Unavailable el)

\(\operatorname{sic} 2 \leftarrow \operatorname{VSR[VBB}+32] \& 0 \times 7 F F F F F F F F F F F F F F F F F F F F F F F F F F F F\)
VSR[VRT+32] \& STC1 | sic2
Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let src2 be the floating-point value in VSR[VRB+32] represented in quad-precision format.
\(\operatorname{src2}\) is placed into VSR[VRT+32] with the sign of src1.
Special Registers Altered:
None
VSR Data Layout for xscpsgnqp
VSR[ VRA+32]
\begin{tabular}{l}
\begin{tabular}{|c|}
\hline SrCl \\
VSR[ VRB +32\(]\) \\
\hline src2 \\
\hline VSR[VRT+32] \\
\hline
\end{tabular} \\
\hline
\end{tabular}

VSX Scalar round \& Convert Double-Precision format to Half-Precision format XX2-form

if MSR.VSX=0 then VSX_Unavailable(l)
reset_flags()
sic \(\leftarrow\) bfp_CONVERT_FROM_BFP64(VSR[BX×32+B]. dword[0])
rnd \(\leftarrow\) bfP_ROUND_TO_BPP16(FPSCR,RN, SIC)
result \(\leftarrow\) bfp_CONVERT_TO_BFP16(rnd)
if(vxsnan_flag) then SetFX(FPSCR. VXSNAN)
if (ox_flag) then SetFXIFPSCR. OX)
if(ux-flag) then SetFX(FPSCR. UX)
if(xx_flag) then SetFX(FPSCR. XX)
vex_flag \(\leftarrow\) FPSCR.VE \& vxsnan_flag
if vex_flag=0 then do
VSR[TXX32+T],hword \([0: 2] \leftarrow 0 \times 0000 \_0000 \_0000\)
VSR[TX×32+T], hword[3] \(\leftarrow\) result \({ }^{-}\)
VSR[TX \(\times 32+T]\). dwor \(d[1] \leftarrow\) OxUUUU_UUUU_UUUU_UUUU
FPSCR.FPRF \(\leftarrow\) fprf_CLASS_BFP16(result)
end
FPSCR.FR \(\leftarrow(\) vex_flag \(=0)\) \& inc_flag

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

If src is an SNaN , the result is the half-precision representation of that SNaN converted to a QNaN .

Otherwise, if \(\operatorname{src}\) is a QNaN , the result is the half-precision representation of that QNaN .

Otherwise, if src is an Infinity, the result is the half-precision representation of Infinity with the same sign as src.

Otherwise, if src is a Zero, the result is the half-precision representation of Zero with the same sign as sic.

Otherwise, the result is the half-precision representation of \(\operatorname{src}\) rounded to half-precision using the rounding mode specified by RN.

The result is zero-extended and placed into doubleword element 0 of VSR[XT].

The contents of doubleword element 1 of VSR[ XT] are undefined.

FPRF is set to the class and sign of the result as represented in half-precision. \(F R\) is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[ XT] and FPRF are not modified, and FR and FI are set to 0 .

\section*{Special Registers Altered:}

FPRF FR FI
FX VXSNan OX UX XX

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

\section*{VSR Data Layout for xscvdphp}


VSX Scalar Convert Double-Precision format to Quad-Precision format X-form

if MSR.VSX=0 then VSX_Unavailablell
src \(\leftarrow\) bfp_CONVERT_FROM_BFP64(VSR[VRB+32]. dword[0])
if src.class. SNaN then
result \(\leftarrow\) bfp_CONVERT_TO_BFP128(bfp_QUIET(src))
else
result \(\leftarrow\) bfp_CONVERT_TO_BFP128(src)
vxsnan_flag \(\leftarrow\) src.class. SNaN
if(vxsnan_flag) then SetFX(FPSCR. VXSNAN)
vex_flag \(\leftarrow\) FPSCR.VE \& vxsnan_flag
if vex_flag=0 then do
VSR[VRT+32] \(\leftarrow\) result
FPSCR.FPRF \(\leftarrow\) fprf_CLASS_BPP128(result)
end
\(F P S C R, F R \leftarrow 0\)
FPSCR,FI \(\leftarrow 0\)
Let \(\operatorname{src}\) be the floating-point value in doubleword element 0 of VSR[VRB+32] represented in double-precision format.

Src is placed into VSR[VRT+32] in quad-precision format.

If \(\operatorname{src}\) is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN is set to 1 .

FPRF is set to the class and sign of the result.
FR is set to \(0 . \mathrm{Fl}\) is set to 0 .
If a trap-enabled Invalid Operation exception occurs, VSR[XT] and FPRF are not modified.

Special Registers Altered:
FPRF FR (set to 0) FI (set to 0)
FX VXSNAN
| VSR Data Layout for xscvdpqp
VSR[ VRB+32]
\(\square\)
VSR[ VRT +32]
tgt

VSX Scalar round Double-Precision to single-precision and Convert to Single-Precision format XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline xscvdp & & XT, XB & & & & & \\
\hline \[
60
\] & \[
{ }_{6} \mathrm{~T}
\] & \[
{ }_{11} \text { III }
\] & & B & 21 & 265 & \(\left|\begin{array}{l}\text { BXTX } \\ 3031\end{array}\right|\) \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { reset_xflags() } \\
& \mathrm{src} \leftarrow \mathrm{VSR}[32 \times B X+B] \text {. dword [0] } \\
& \text { result } \leftarrow \text { ConvertDPtoSP(src) } \\
& \text { if(vxsnan_flag) then SetFX(FPSCR.VXSNAN) } \\
& \text { if(xx_flag) then SetFX(FPSCR.XX) } \\
& \text { if(ox_flag) then SetFX(FPSCR.OX) } \\
& \text { if(ux_flag) then SetFX(FPSCR.UX) } \\
& \text { vex_flag } \leftarrow \text { FPSCR.VE \& vxsnan_flag } \\
& \text { if( ~vex_flag) then do } \\
& \text { VSR }[32 \times T X+T] \text {.word }[0] \leftarrow \text { result } \\
& \text { VSR }[32 \times T X+T] \text {. Word }[1] \leftarrow \text { exUUUU_UUUU } \\
& \text { VSR }[32 \times T X+T] \text {. Word }[2] \leftarrow \text { exUUUU_UUUU } \\
& \text { VSR[32xTX+T]. word [3] \& 0xUUUU_UUUU } \\
& \text { FPSCR.FPRF } \leftarrow \text { ClassSP(result) } \\
& \text { FPSCR.FR } \leftarrow \text { inc_flag } \\
& \text { FPSCR.FI } \leftarrow \text { xx_flag } \\
& \text { end } \\
& \text { else do } \\
& \text { FPSCR.FR } \leftarrow 0 \text { b0 } \\
& \text { FPSCR.FI } \leftarrow 0 \text { b0 }
\end{aligned}
\]

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(\mathrm{src}^{\circ}\) be the double-precision floating-point value in doubleword element 0 of VSR[XB].

If \(s r^{c}\) is a SNaN , the result is src converted to a QNaN (i.e., bit 12 of src is set to 1 ). VXSNAN is set to 1 .

Otherwise, if \(\operatorname{src}\) is a QNaN, an Infinity, or a Zero, the result is src .

Otherwise, the result is sic rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into word element 0 of VSR[ XT] in single-precision format.

The contents of word elements 1,2 , and 3 of VSR[ XT] are undefined.

FPRF is set to the class and sign of the result. \(F R\) is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

Special Registers Altered
FPRF FR FI FX OX UX XX VXSNAN

\section*{VSR Data Layout for xscvdpsp}

STC = VSR[XB]
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\operatorname{VSR}[X T]\)
\begin{tabular}{|l|l|ll|}
\hline SP & undefined & undefined \\
\hline 0 & 32 & 64 & 127 \\
\hline
\end{tabular}

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

\section*{VSX Scalar Convert Scalar Single-Precision to Vector Single-Precision format Non-signalling XX2-form}
\[
\text { xscvdpspn } \quad \mathrm{XT}, \mathrm{XB}
\]


> reset_xflags()
> src \(\leftarrow\) VSR[32×BX \(+B]\). dword[0]
> result \(\leftarrow\) ConvertDPtoSP_NS(src)
> VSR \([32 \times T X+T]\) word \([0] \leftarrow\) result
> VSR[32xTX+T]. word[1] \& 0xUUUU_UUUU
> VSR \([32 \times T X+T]\).word \([2] \leftarrow\) 0xUUUU_UUUU
> VSR \([32 x T X+T]\).word \([3] \leftarrow 0 x U U U U \_U U U U\)

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(\operatorname{src}\) be the single-precision floating-point value in doubleword element 0 of VSR[XB] represented in double-precision format.
\(\operatorname{src}\) is placed into word element 0 of VSR[XT] in single-precision format.

The contents of word elements 1,2 , and 3 of VSR[ XT] are undefined.

Special Registers Altered
None
VSR Data Layout for xscvdpspn
SrC = VSR[XB]
\begin{tabular}{|c|c|}
\hline\(S P\) & unused \\
\hline
\end{tabular}
tgt \(=\) VSR[ XT]
\begin{tabular}{|l|l|l|l|}
\hline & SP & undefined & undefined \\
undefined \\
\hline 0 & 32 & 64 & 96
\end{tabular}

\section*{Programming Note}
xscvdpsp should be used to convert a scalar double-precision value to vector single-precision format.
xscvdpspn should be used to convert a scalar single-precision value to vector single-precision format.

\section*{VSX Scalar truncate Double-Precision to integer and Convert to Signed Integer Doubleword format with Saturate XX2-form}
xscvdpsxds \(\quad\) XT,XB

\[
\begin{aligned}
& X T \quad \leftarrow T X \| T \\
& X B \quad \leftarrow B X|\mid B \\
& \text { reset_xflags() } \\
& \text { result }\{0: 63\} \leftarrow \text { ConvertDPtoSD (VSR }[X B]\{0: 63\}) \\
& \text { if(vxsnan_flag) then SetFX(VXSNAN) } \\
& \text { if(vxcvi_flag) then SetFX(vXCVI) } \\
& \text { if(xx_flag) then } \operatorname{SetFX}(X X) \\
& \text { vex_flag } \leftarrow \text { VE \& (vxsnan_flag | vxcvi_flag) } \\
& \text { if( ~vex_flag ) then do } \\
& \text { VSR[XT] \& result || 0xUUUU_UUUU_UUUU_UUUU } \\
& \text { FPRF } \leftarrow \text { ObUUUUU } \\
& \mathrm{FR} \quad \leftarrow \text { inc_flag } \\
& \text { FI } \leftarrow \text { xx_flag } \\
& \text { end } \\
& \text { else do } \\
& \mathrm{FR} \leftarrow 0 \mathrm{~b} 0 \\
& \mathrm{FI} \leftarrow 0 \mathrm{bo} \\
& \text { end }
\end{aligned}
\]

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(\operatorname{src}\) be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

If src is a NaN , the result is the value \(0 \times 8000 \_0000 \_0000 \_0000\) and VXCVI is set to 1 . If \(\operatorname{sic}\) is an SNaN, VXSNAN is also set to 1 .

Otherwise, src is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than \(2^{63} \cdot 1\), the result is OXTFFF_FFFF_FFFF_FFF and VXCVI is set to 1 .

Otherwise, if the rounded value is less than \(\cdot 2^{63}\), the result is \(0 \times 80000_{-} 0000_{-} 00000_{-} 0000\) and VXCVI is set to 1 .

Otherwise, the result is the rounded value converted to 64-bit signed-integer format, and if the result is inexact (i.e., not equal to src), \(X X\) is set to 1 .

If a trap-enabled invalid operation exception occurs,
- VSR[ XT] and FPRF are not modified
- \(F R\) and \(F I\) are set to 0 .

\section*{Version 3.0}

Otherwise,
- The result is placed into doubleword element 0 of VSR \([X T]\). The contents of doubleword element 1 of VSR[ \(X T]\) are undefined.
- FPRF is set to an undefined value.
- \(F R\) is set to indicate if the result was incremented when rounded.
- Fl is set to indicate the result is inexact.

See Table 66.

\section*{Special Registers Altered}
```

FPRF=ObUUUUU FR FI FX XX VXSNAN VXCVI

```

\section*{VSR Data Layout for xscvdpsxds}
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline SD & \multicolumn{2}{c|}{ undefined } \\
\hline 0 & 64
\end{tabular}

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

\section*{- Programming Note}
xscvdpsxds rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Double-Precision Integer instruction that corresponds to the desired rounding mode, including xsrdpic which uses the rounding mode specified by RN.
\begin{tabular}{|c|c|c|c|c|}
\hline & 山 & 山 &  & Returned Results and Status Setting \\
\hline \multirow[t]{2}{*}{SrC \(\leq\) Nmin-1} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\) \\
\hline & 1 & - & - & \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\), error() \\
\hline \multirow[t]{2}{*}{Nmin-1 < SrC < Nmin} & \multirow[t]{2}{*}{-} & 0 & yes & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\) \\
\hline & & 1 & yes & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\), error() \\
\hline SrC \(=\) Nmin & - & - & no & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0\) \\
\hline \multirow{3}{*}{Nmin < SrC < Nmax} & \multirow{3}{*}{-} & - & no & T(ConvertDPtoSD(RoundToDPintegerTrunc(src)) , FR \(\leftarrow 0, \mathrm{Fl} \leftarrow 0\) \\
\hline & & 0 & yes & T (ConvertDPtoSD(RoundToDPintegerTrunc(src))), \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\) \\
\hline & & 1 & yes & T (ConvertDPtoSD(RoundToDPintegerTrunc(src))), FR\&0, Fl¢1, fx(XX), error() \\
\hline SrC \(=\) Nmax & - & - & no & \begin{tabular}{l}
\(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{~F} \leftarrow 0\) \\
Note: This case cannot occur as Nmax is not representable in DP format but is included here for completeness.
\end{tabular} \\
\hline \multirow[b]{2}{*}{Nmax < src < Nmax+1} & \multirow[t]{2}{*}{-} & 0 & yes & \(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\) \\
\hline & & 1 & yes & \(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\), error() \\
\hline \multirow[t]{2}{*}{src \(\geq\) Nmax +1} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\) \\
\hline & 1 & - & - & \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}\) (VXCVI), error() \\
\hline \multirow[t]{2}{*}{src is a QNaN} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\) \\
\hline & 1 & - & - & \(F R \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\), error() \\
\hline \multirow[t]{2}{*}{src is a SNaN} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI}), \mathrm{fx}(\mathrm{VXSNAN})\) \\
\hline & 1 & - & - & \(F R \leftarrow 0, F l \leftarrow 0, f x(V X C V I), ~ f x(V X S N A N), ~ e r r o r() ~\) \\
\hline \multicolumn{5}{|l|}{Explanation:} \\
\hline \multicolumn{5}{|c|}{\(F X\) is set to 1 if \(\mathrm{x}=0 . \mathrm{x}\) is set to 1 .} \\
\hline \multicolumn{5}{|r|}{The system error handler is invoked for the trap-enabled exception if the FE0 and FE1 bits in the Machine State Register are set to any mode other than the ignore-exception mode.} \\
\hline \multicolumn{5}{|r|}{The smallest signed integer doubleword value, -2 \({ }^{63}\) (0x8000_0000_0000_0000).} \\
\hline \multicolumn{5}{|r|}{The largest signed integer doubleword value, \(2^{63}-1\) ( \(0 \times 7\) FFF_FFFF_FFFF_FFFF).} \\
\hline \multicolumn{5}{|r|}{The double-precision floating-point value in doubleword element 0 of VSR[XB].} \\
\hline \multicolumn{5}{|r|}{The signed integer doubleword value \(x\) is placed in doubleword element 0 of VSR[XT]. The contents of doubleword element 1 of VSR[XT] are undefined.} \\
\hline
\end{tabular}

Table 66.Actions for xscvdpsxds

\section*{VSX Scalar truncate Double-Precision to integer and Convert to Signed Integer Word format with Saturate XX2-form}
\[
\text { xscvdpsxws } \quad \text { XT,XB }
\]
\begin{tabular}{|l|ll|l|l|l|l|l|}
\hline 60 & & T & & & III & & B \\
\hline 0 & & 6 & & & & 88 & \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline XT & \(\leftarrow T X \| T\) \\
\hline XB & \(\leftarrow B X|\mid B\) \\
\hline inc_flag & \(\leftarrow 0 \mathrm{bo}\) \\
\hline \multicolumn{2}{|l|}{reset_xflags()} \\
\hline \multicolumn{2}{|l|}{result \(\{0: 31\} \leqslant\) ConvertDPtoSW(VSR[XB]\{0:63\})} \\
\hline \multicolumn{2}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{2}{|l|}{if(vxcvi_flag) then SetFX(VXCVI)} \\
\hline \multicolumn{2}{|l|}{if(xx_flag) then SetFX(XX)} \\
\hline vex_flag & \(\leftarrow\) VE \& (vxsnan_flag | vxcvi_flag) \\
\hline \multicolumn{2}{|l|}{if( ~vex_flag ) then do} \\
\hline \multicolumn{2}{|l|}{VSR[XT] ¢ 0xUUUU_UUUU || result || 0xUUUU_UUUU_UUUU_UUUU} \\
\hline \multicolumn{2}{|l|}{FPRF \(\leftarrow\) ObUUUUU} \\
\hline \multicolumn{2}{|l|}{FR \(\leftarrow\) inc_flag} \\
\hline \multicolumn{2}{|l|}{FI \(\leftarrow\) xx_flag} \\
\hline \multicolumn{2}{|l|}{end} \\
\hline \multicolumn{2}{|l|}{else do} \\
\hline FR & \(\leftarrow 0 \mathrm{bo}\) \\
\hline FI & \(\leftarrow 0 \mathrm{bo}\) \\
\hline end & \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(\operatorname{src}\) be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

If src is a NaN , the result is the value \(0 \times 8000 \_0000\) and VXCVI is set to 1 . If src is an SNaN, VXSNAN is also set to 1.

Otherwise, src is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than \(2^{31} .1\), the result is \(0 \times 7 F F F\) FFFF and \(V X C V I\) is set to 1 .

Otherwise, if the rounded value is less than \(\cdot 2^{31}\), the result is \(0 \times 8000 \_0000\) and VXCVI is set to 1 .

Otherwise, the result is the rounded value converted to 32-bit signed-integer format, and if the result is inexact (i.e., not equal to src ), \(X X\) is set to 1 .

If a trap-enabled invalid operation exception occurs,
- VSR[XT] and FPRF are not modified
- \(F R\) and \(F I\) are set to 0 .

Otherwise,
- The result is placed into word element 1 of VSR[XT]. The contents of word elements 0,2 , and 3 of VSR[ XT] are undefined.
- FPRF is set to an undefined value.
- \(F R\) is set to indicate if the result was incremented when rounded.
- Fl is set to indicate the result is inexact.

See Table 67.

\section*{Special Registers Altered}

FPRF=ObUUUUU FR FI FX XX VXSNAN VXCVI

\section*{VSR Data Layout for xscvdpsxws}
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|l|l|}
\hline undefined & & SW & \\
\hline & undefined & \\
\hline 0 & 32 & 64 &
\end{tabular}

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

\section*{- Programming Note}
xscvdpsxws rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Double-Precision Integer instruction that corresponds to the desired rounding mode, including xsrdpic which uses the rounding mode specified by RN.


Table 67.Actions for xscvdpsxws

VSX Scalar truncate Double-Precision integer and Convert to Unsigned Integer Doubleword format with Saturate XX2-form
xscvdpuxds XT,XB



Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(\operatorname{src}\) be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

If \(\operatorname{src}\) is a NaN, the result is the value \(0 \times 0000-0000,0000 \_0000\) and VXCVI is set to 1 . If \(\operatorname{src}\) is an SNaN, VX'SNAN is also set to 1 .

Otherwise, \(\operatorname{src}\) is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than \(2^{64} \cdot 1\), the result is


Otherwise, if the rounded value is less than 0 , the result is \(0 \times 00000_{-0000} 00000_{-} 0000\) and VXCVI is set to 1 .

Otherwise, the result is the rounded value converted to 64 -bit unsigned-integer format, and if the result is inexact (i.e., not equal to \(s r^{c}\) ), \(X X\) is set to 1 .

If a trap-enabled invalid operation exception occurs,
- VSR[XT] and FPRF are not modified
- \(F R\) and \(F I\) are set to 0 .

Otherwise,
- The result is placed into doubleword element 0 of VSR[ XT]. The contents of doubleword element 1 of VSR[XT] are undefined.
- FPRF is set to an undefined value.
- \(F R\) is set to indicate if the result was incremented when rounded.
- Fl is set to indicate the result is inexact.

See Table 68.

\section*{Special Registers Altered}
\[
F P R F=O b U U U U U \quad F R \quad F I \quad F X \quad X X \text { VXSNAN VXCVI }
\]

\section*{VSR Data Layout for xscvdpuxds}
\[
\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]
\]
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline & \multicolumn{2}{c|}{ undefined } \\
\hline 0 & 64 \\
\hline
\end{tabular}

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

\section*{- Programming Note}
xscvdpuxds rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Double-Precision Integer instruction that corresponds to the desired rounding mode, including xsrdpic which uses the rounding mode specified by RN.
\begin{tabular}{|c|c|c|c|c|}
\hline & ш & 山 &  & Returned Results and Status Setting \\
\hline \multirow[b]{2}{*}{Src \(\leq\) Nmin-1} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\) \\
\hline & 1 & - & - & \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\), error() \\
\hline \multirow[t]{2}{*}{Nmin-1 < SrC < Nmin} & \multirow[t]{2}{*}{-} & 0 & yes & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\) \\
\hline & & 1 & yes & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\), error() \\
\hline src \(=\) Nmin & - & - & no & T (Nmin), \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0\) \\
\hline \multirow{3}{*}{Nmin < SrC < Nmax} & \multirow{3}{*}{-} & - & no & T (ConvertDPtoUD(RoundToDPintegerTrunc(src)) ), FR \(\leftarrow 0, \mathrm{Fl} \leftarrow 0\) \\
\hline & & 0 & yes & T (ConvertDPtoUD(RoundToDPintegerTrunc(src))), \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\) \\
\hline & & 1 & yes & T (ConvertDPtoUD(RoundToDPintegerTrunc(src))), FR\&0, Fl<1, fx(XX), error() \\
\hline SrC \(=\) Nmax & - & - & no & \begin{tabular}{l}
\[
\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0
\] \\
Note: This case cannot occur as Nmax is not representable in DP format but is included here for completeness.
\end{tabular} \\
\hline \multirow[t]{2}{*}{Nmax < SrC < Nmax+1} & \multirow[t]{2}{*}{-} & 0 & yes & \(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\) \\
\hline & & 1 & yes & \(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\), error () \\
\hline \multirow[b]{2}{*}{\(\operatorname{src} \geq \mathrm{Nmax}+1\)} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\) \\
\hline & 1 & - & - & \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\), error() \\
\hline \multirow[b]{2}{*}{src is a QNaN} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\) \\
\hline & 1 & - & - & \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\), error() \\
\hline \multirow[b]{2}{*}{src is a SNaN} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI}), \mathrm{fx}(\mathrm{VXSNAN})\) \\
\hline & 1 & - & - & \(F R \leftarrow 0, F l \leftarrow 0, f x(V X C V I), ~ f x(V X S N A N), ~ e r r o r() ~\) \\
\hline \multicolumn{5}{|l|}{Explanation:} \\
\hline \multicolumn{5}{|c|}{\(F X\) is set to 1 if \(\mathrm{x}=0 . \mathrm{x}\) is set to 1 .} \\
\hline \multicolumn{5}{|r|}{The system error handler is invoked for the trap-enabled exception if the FE0 and FE1 bits in the Machine State Register are set to any mode other than the ignore-exception mode.} \\
\hline \multicolumn{5}{|r|}{The smallest unsigned integer doubleword value, 0 (0x0000_0000_0000_0000).} \\
\hline \multicolumn{5}{|r|}{The largest unsigned integer doubleword value, \(2^{64}-1\) (0xFFFF_FFFF_FFFF_FFFF).} \\
\hline \multicolumn{5}{|r|}{The double-precision floating-point value in doubleword element 0 of VSR[XB].} \\
\hline \multicolumn{5}{|r|}{The unsigned integer doubleword value x is placed in doubleword element 0 of VSR[XT]. The contents of doubleword element 1 of VSR[XT] are undefined.} \\
\hline
\end{tabular}

Table 68.Actions for xscvdpuxds

VSX Scalar truncate Double-Precision to integer and Convert to Unsigned Integer Word format with Saturate XX2-form
\[
\text { xscvdpuxws } \quad \text { XT,XB }
\]
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 60 & & T & \multicolumn{1}{c|}{ III } & & B & & 72 \\
\hline 0 & & 6 & & 11 & & 16 & \\
21 & & \(307 X\) \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline XT & \(\leftarrow T X \| T\) \\
\hline XB & \(\leftarrow B X|\mid B\) \\
\hline inc_flag & \(\leftarrow 0 \mathrm{bO}\) \\
\hline \multicolumn{2}{|l|}{reset_xflags()} \\
\hline \multicolumn{2}{|l|}{result \(\{0: 31\} \leftarrow\) ConvertDPtoUW(VSR[XB] \(30: 63\}\) )} \\
\hline \multicolumn{2}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{2}{|l|}{if(vxcvi_flag) then SetFX(vXCVI)} \\
\hline \multicolumn{2}{|l|}{if(xx_flag) then \(\operatorname{SetFX}(\mathrm{XX})\)} \\
\hline vex_flag & \(\leftarrow\) VE \& (vxsnan_flag | vxcvi_flag) \\
\hline \multicolumn{2}{|l|}{if( ~vex_flag ) then do} \\
\hline \multicolumn{2}{|l|}{VSR[XT] \& 0xUUUU_UUUU || result || 0xUUUU_UUUU_UUUU_UUUU} \\
\hline \multicolumn{2}{|l|}{FPRF \(\leftarrow\) ObUUUUU} \\
\hline \multicolumn{2}{|l|}{FR \(\leftarrow\) inc_flag} \\
\hline \multicolumn{2}{|l|}{FI \(\leftarrow\) xx_flag} \\
\hline \multicolumn{2}{|l|}{end} \\
\hline \multicolumn{2}{|l|}{else do} \\
\hline FR & \(\leftarrow 0 \mathrm{bo}\) \\
\hline FI & \(\leftarrow 0 \mathrm{~b} 0\) \\
\hline end & \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(\operatorname{src}\) be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

If src is a NaN , the result is the value \(0 \times 0000 \_0000\) and VXCVI is set to 1 . If src is an SNaN , VXSNAN is also set to 1.

Otherwise, \(\operatorname{src}\) is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than \(2^{32} \cdot 1\), the result is \(0 \times F F F F \_F F F\) and \(V X C V I\) is set to 1 .

Otherwise, if the rounded value is less than 0 , the result is \(0 \times 0000 \_0000\) and VXCVI is set to 1 .

Otherwise, the result is the rounded value converted to 32-bit unsigned-integer format, and if the result is inexact (i.e., not equal to \(s r^{c}\) ), \(X X\) is set to 1 .

If a trap-enabled invalid operation exception occurs,
- VSR[XT] and FPRF are not modified
- \(F R\) and \(F I\) are set to 0 .

Otherwise,
- The result is placed into word element 1 of VSR[XT]. The contents of word elements 0,2 , and 3 of \(\operatorname{VSR}[X T]\) are undefined.
- FPRF is set to an undefined value.
- \(F R\) is set to indicate if the result was incremented when rounded.
- Fl is set to indicate the result is inexact.

See Table 69.

\section*{Special Registers Altered}
\[
F P R F=O b U U U U U \quad F R \quad F I \quad F X \quad X X \text { VXSNAN VXCVI }
\]

\section*{VSR Data Layout for xscvdpuxws}
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|ll|}
\hline undefined & UW & \multicolumn{2}{c|}{ undefined } \\
\hline 0 & 32 & 64 & 127 \\
\hline
\end{tabular}

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

\section*{Programming Note}
xscvdpuxws rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Double-Precision Integer instruction that corresponds to the desired rounding mode, including xsrdpic which uses the rounding mode specified by RN.
\begin{tabular}{|c|c|c|c|c|}
\hline & 山 & 岗 & Inexact? ( RoundToDPintegerTrunc(src) \(=\mathrm{src}\) ) &  \\
\hline \multirow[b]{2}{*}{Src \(\leq\) Nmin-1} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\) \\
\hline & 1 & - & - & FR↔0, Fl¢ \(\leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\), error() \\
\hline \multirow[b]{2}{*}{Nmin-1 < SrC < Nmin} & \multirow{2}{*}{-} & 0 & yes & T (Nmin), \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\) \\
\hline & & 1 & yes & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\), error() \\
\hline SrC \(=\) Nmin & - & - & no & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0\) \\
\hline \multirow{3}{*}{Nmin < SrC < Nmax} & \multirow{3}{*}{-} & - & no & T (ConvertDPtoUW(RoundToDPintegerTrunc(src))), FR \(\leftarrow 0, \mathrm{Fl} \leftarrow 0\) \\
\hline & & 0 & yes & T (ConvertDPtoUW(RoundToDPintegerTrunc(src))), \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\) \\
\hline & & 1 & yes & T (ConvertDPtoUW(RoundToDPintegerTrunc(src))), FR \(\leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\), error() \\
\hline SrC \(=\) Nmax & - & - & no & \(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0\) \\
\hline \multirow[t]{2}{*}{Nmax < Src < Nmax+1} & \multirow[t]{2}{*}{-} & 0 & yes & \(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\) \\
\hline & & 1 & yes & \(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\), error() \\
\hline \multirow[b]{2}{*}{Src \(\geq\) Nmax +1} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmax}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\) \\
\hline & 1 & - & - & FR↔0, Fl \(\leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\), error() \\
\hline \multirow[b]{2}{*}{src is a QNaN} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\) \\
\hline & 1 & - & - & \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI})\), error() \\
\hline \multirow[t]{2}{*}{src is a SNaN} & 0 & - & - & \(\mathrm{T}(\mathrm{Nmin}), \mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI}), \mathrm{fx}(\mathrm{VXSNAN})\) \\
\hline & 1 & - & - & \(\mathrm{FR} \leftarrow 0, \mathrm{Fl} \leftarrow 0, \mathrm{fx}(\mathrm{VXCVI}), \mathrm{fx}(\mathrm{VXSNAN})\), error() \\
\hline \multicolumn{5}{|l|}{Explanation:} \\
\hline \multicolumn{5}{|c|}{FX is set to 1 if \(\mathrm{x}=0 . \mathrm{x}\) is set to 1.} \\
\hline \multicolumn{5}{|l|}{The system error handler is invoked for the trap-enabled exception if the FEO and FE1 bits in the Machine State Register are set to any mode other than the ignore-exception mode.} \\
\hline \multicolumn{5}{|r|}{The smallest unsigned integer word value, 0 (0x0000_0000).} \\
\hline \multicolumn{5}{|r|}{The largest unsigned integer word value, \(2^{32}-1\) (0xFFFF_FFFF).} \\
\hline \multicolumn{5}{|r|}{The double-precision floating-point value in doubleword element 0 of VSR[XB].} \\
\hline \multicolumn{5}{|r|}{The unsigned integer word value \(x\) is placed in word element 1 of VSR[ \(X T]\). The contents of word elements 0,2 , and 3 of VSR[XT] are undefined.} \\
\hline
\end{tabular}

Table 69.Actions for xscvdpuxws

VSX Scalar Convert Half-Precision format to Double-Precision format XX2-form
xscvhpdp XT,XB

if MSR.VSX=0 then VSX_Unavailablell
reset_flags()
src \(\leftarrow\) bfp_CONVERT_FROM_BFP16(VSR[BX×32+B], hword[3])
if src.class. \(5 N a N=1\) then
result \(\leftarrow\) bfp_CONVERT_TO_BFP64(bfp_QUIET(sic))
else
result \(\leftarrow\) bfp_CONVERT_TO_BFP64(src)
vxsnan_flag \(\leftarrow\) src.class. SNaN
if(vxsnan_flag) then SetFX(FPSCR. VXSNAN)
vex_flag \(\leftarrow\) FPSCR.VE \& vxsnan_flag
if vex_flag=0 then do
\[
\begin{aligned}
& \operatorname{VSR}[T X \times 32+T] \text {.dword }[0] \leftarrow \text { result } \\
& \text { VSR[TX×32 +T].dword[1] © OXUUUU_UUUU_UUUU_UUUU } \\
& \text { FPSCR.FPRF } \leftarrow \text { fprf_CLASS_BFP64(result) } \\
& \text { end } \\
& \text { FPSCR.FR } \leftarrow 0 \\
& \text { FPSCR.FI } \leftarrow 0
\end{aligned}
\]

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the half-precision floating-point value in the rightmost halfword of doubleword element 0 of VSR[ XB].

If \(s r c\) is an SNaN , the result is the double-precision representation of that SNaN converted to a QNaN .

Otherwise, if \(\operatorname{src}\) is a QNaN , the result is the double-precision representation of that QNaN .

Otherwise, if src is an Infinity, the result is the double-precision representation of Infinity with the same sign as src.

Otherwise, if \(\operatorname{src}\) is a Zero, the result is the double-precision representation of Zero with the same sign as src.

Otherwise, if \(s r c\) is a denormal value, the result is the normalized double-precision representation of src .

Otherwise, the result is the double-precision representation of sic .

The result is placed into doubleword element 0 of VSR[ XT].

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result as represented in half-precision.

If a trap-enabled invalid operation exception occurs, VSR[ XT] and FPRF are not modified.
\(F R\) is set to \(0 . F \mid\) is set to 0 .

\section*{Special Registers Altered:}

FPRF FR (set to 0) FI (set to 0)
FX VXSNAN

VSR Data Layout for xscvhpdp


\section*{VSX Scalar round \& Convert Quad-Precision format to Double-Precision format [using round to Odd] X-form}
\begin{tabular}{lll} 
xscvqpdp & VRT,VRB & \((R O=0)\) \\
xscvqpdpo & VRT,VRB & \((R 0=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 63 & \({ }_{6}\) VRT & \multicolumn{1}{c|}{20} & VRB & & 836 & \begin{tabular}{l} 
R0 \\
31
\end{tabular} \\
\hline
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable()
reset_xflags()
src \&bfp CONVERT FROM BFP128(VSR[ VBB+32])
Ind \&bfp_ROUNDTÖ_bFP6̈4(RO,FPSCR.RN,STC)
result \&bfp_CONVERT_TO_BFP64(rnd)
if(vxsnan_flag) then SetFX(FPSCR.vXSNaN)
if(oxflag) then SetFX(FPSCR.OX)
if(ux_flag) then SetFX(FPSCR.UX)
if(x__flag) then SetFX(FPSCR. XX)
vex_flag\&FPSCR.VE\&vxsnan_flag
if vex flag=0 then do
VSR[VRT+32].dword[0] \leftarrow result
VSR[VRT+32], dwor dl [1] \& Ox0000_0000_0000_0000
FPSCR.FPRF}\leftarrow\mathrm{ forf_CLASS_BPP64(result)
end
FPSCR.FR\&(vxsnan_flag=0)\& inc flag
FPSCR.FI\leftarrow(vxsnan_flag=0)\& \&x_flag

```

Let src be the quad-precision floating-point value in VSR[ VRB 32 ].

If src is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN is set to 1 .

If src is a Signalling NaN , the result is the Quiet NaN corresponding to the Signalling NaN , with the significand truncated to the rounding precision.

Otherwise, if src is a Quiet NaN , then the result is src with the significand truncated to double-precision.

Otherwise, if src is an Infinity or a Zero, the result is src.

Otherwise, do the following.
If src is Tiny (i.e., the unbiased exponent is less than -1022) and \(U E=0\), the significand is shifted right \(N\) bits, where \(N\) is the difference between -1022 and the unbiased exponent of src. The exponent of \(s r^{c}\) is set to the value -1022 .

If \(R O=1\), let the rounding mode be Round to Odd. Otherwise, let the rounding mode be specified by RN. Unless the result is an Infinity or a Zero, the intermediate result is rounded to double-precision (i.e., 11-bit exponent range and 53-bit significand precision) using the specified rounding mode.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[VRT+32] in double-precision format. The contents of doubleword element 1 of VSR[ VRT +32 ] are set to 0.

FPRF is set to the class and sign of the result as represented in double-precision format. \(F R\) is set to indicate if the rounded result was incremented. Fl is set to indicate the result is inexact.

If a trap-disabled Invalid Operation exception occurs,
\(F R\) and Fl are set to 0 .
If a trap-enabled Invalid Operation exception occurs, VSR[VRT+32] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered:}

FPRF FR FI
FX VXSNAN OX UX XX
| VSR Data Layout for xscvqpdp[o]
VSR[ VRB+32]
\begin{tabular}{l}
\hline \multicolumn{2}{|c|}{ SIC } \\
\begin{tabular}{|c|c|}
\hline VSR[VRT+32] \\
\hline tgt.dword[0] & \(0 \times 0000 \_0000 \_0000 \_0000\) \\
\hline
\end{tabular}
\end{tabular}

\section*{VSX Scalar truncate \& Convert Quad-Precision format to Signed Doubleword format \(X\)-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|l|}{xscvqpsdz VRT,VRB} \\
\hline \[
0
\] & \[
\left.\right|_{6} \text { VRT }
\] & \[
25
\] & \[
{ }_{16} \text { VRB }
\] & 21 & 836 & \begin{tabular}{|r|}
1 \\
31
\end{tabular} \\
\hline
\end{tabular}
if MSR. VSX=0 then VSX_Unavailablell
reset_xflags()
sic \&bf p_CONERT_ _ROM_BFP128 (VSR[VRB+32])
if src.class. QNaN | src.class. SNaN then do
result \(\leftarrow 0 \times 8000\) _0000_0000_0000
vxsnan_flag sicic.class. SÑaN
vxcui_flag \(\leftarrow 1\)
end
else if src.class.Infinity then do
vxcvi_flag \(\leftarrow 1\)
if sicisign \(=0\) then
result \(\leftarrow\) Ox7FFF_FFFF_FFFF_FFF
else
result \(\leftarrow 0 \times 8000\) _0000_0000_0000
end
else if src.class. Zero then
result \(\leftarrow\) Ox0000_0000_0000_0000
else do
ind \& bf p_ROUND_to_ I NTEGER(ObOO1,sic)
if bf cocompare GT(rnd, \(+2^{63} \cdot 1\) ) then do result \(\leftarrow\) Ox 7 FFF_fFF_fFF_fFFF vxcvi_flag \(\leftarrow 1\)
end
else if bfo_compare_lt(ind, \(\cdot 2^{63}\) ) then do result \(\leftarrow 0 \times 8000\) _0000_0000_0000 vxcvi_flag \(\leftarrow 1\)
end
else do result \(\leftarrow\) bf p_CoNvert_TO_S164(rnd) if(xx_flag) then Set \(\mathrm{PX}(\) FPSCR. XX)
end
end
if(uxsnan_flag) then Setfx(fpscr.vxsuan)
if(uxcvi_flag) then SetFX(FPSCR. VXCVI)
vx_flag < vxsnan_flag \| vxcvi_flag
ex_flag \(\leftarrow\) FPSCR.VE \(\& V x_{-} f \mid a g\)
if ex_flag=0 then do
VSR[VRT+32]. dword[0] \(\leftarrow\) result
VSR[VRT+32].dword [1] \(\leftarrow\) Ox0000_0000_0000_0000
end
FPSCR. \(F R \leftarrow\left(v x \_f \mid a g=0\right) \& i n c \_f \mid a g\)
FPSCR.FI \(\leftarrow\left(v x \_f \mid a g=0\right) \& x x_{ـ} f \mid a g\)
Let \(\operatorname{src}\) be the quad-precision floating-point value in VSR[ VRB +32 ].

If src is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN and VXCVI are set to 1.

If src is a Quiet NaN or an Infinity, an Invalid Operation exception occurs and VXCVI is set to 1 .

If \(5 r c\) is a NaN , the result is \(0 \times 80000_{-} 0000_{\mathrm{Z}} 0000 \_0000\).
Otherwise, if src is a Zero, the result is \(0 \times 0000 \_0000 \_0000 \_0000\).

Otherwise, if src is +Infinity, the result is OXTFFF_FFF_FFF_FFFF.

Otherwise, if src is - Infinity, the result is \(0 \times 8000 \_0000 \_0000 \_0000\).

Otherwise, do the following.
Let rnd be the value src truncated to a floating-point integer.

If \(r n d\) is greater than \(+2^{63}\). 1 , an Invalid Operation exception occurs, \(V X C V I\) is set to 1 , and the result is \(\mathrm{OXPFFF}_{\ldots}\) FFFF_FFF_FFFF.

Otherwise, if rnd is less than \(\cdot 2^{63}\), an Invalid Operation exception occurs, VXCVI is set to 1 , and the result is \(0 \times 8000 \_0000 \_0000 \_0000\).

Otherwise, the result is the value rnd, and an Inexact exception occurs if \(r\) nd is inexact (i.e., \(r\) nd is not equal to src ).

The result is placed into doubleword element 0 of VSR[ VRT +32 ] in signed integer format.

The contents of doubleword element 1 of VSR[VRT+32] are set to 0 .

FPRF is set to undefined. FR is set to 0 . Fl is set to indicate if the rounded result is inexact.

If an Invalid Operation exception occurs, FR and FI are set to 0 .

If a trap-enabled Invalid Operation exception occurs, VSR[ VRT +32 ] and FPRF are not modified.

See Table 66, "Actions for xscvdpsxds," on page 541.
Special Registers Altered:
FPRF (undefined) FR FI FX VXSNAN VXCVI XX
I VSR Data Layout for xscvqpsdz
I VSR[VRB+32]
I \(\square\)
VSR[ VRT+32]
\begin{tabular}{|c|c|}
\hline tgt.dword[0] & \(0 \times 0000 \_0000 \_0000 \_0000\) \\
\hline
\end{tabular}

\| Table 70. Actions for xscvqpsdz

\section*{VSX Scalar truncate \& Convert Quad-Precision format to Signed Word format X-form}
xscyqpswz VRT,VRB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 63 & VRT & 9 & VRB & & 836 & 1 \\
\hline 0 & & 6 & 11 & & 16 & 21 \\
\hline
\end{tabular}
if MSR. VSX=0 then VSX_Unavailablell
reset_xflags()
sic \&bf p_CONERT_ _ROM_BFP128 (VSR[VRB+32])
if src.class. QNaN | src.class. SNaN then do
result \(\leftarrow\) OxFFFFFFFF_8000_0000
vxsnan_flag sicc.class. SÑaN
vxcui_flag \(\leftarrow 1\)
end
else if src.class.Infinity then do
vxcri_flag \(\leftarrow 1\)
if srcisign \(=0\) then
result \(\leftarrow 0 \times 0000\) _0000_7FFF_FFF
else
result \(\leftarrow\) OxFFFF_fFFF_8000_0000
end
else if src.class.Zero then
result \(\leftarrow\) Ox0000_0000_0000_0000
else do
ind \& bf p_ROUND_to_ I NTEGER(ObOO1,sic)
if bf coCOMPARE GT(rnd, \(+2^{31} \cdot 1\) ) then do result \(\leftarrow 0 x 0000\) _0000_7FFF_FFF vxcvi_flag \(\leftarrow 1\)
end
else if bfp_compare_Lt(rnd, \(\cdot 2^{31}\) ) then do result \(\leftarrow\) OxFFFF_FFFF_8000_0000 vxcvi_flag \(\leftarrow 1\)
end
else do result \(\leftarrow\) bf p_CoNvert_TO_S164(rnd) if(xx_flag) then SetFX(FPSCR. XX)
end
end
if(uxsnan_flag) then Setfx(fpscr.vxsuan)
if(uxcvi_flag) then Set FX(FPSCR. VXCVI)
vx_flag < vxsnan_flag \| vxcvi_flag
ex_flag \(\leftarrow\) FPSCR.VE \(\& v x\) flag
if ex_flag=0 then do
VSR[ VRT +32\(]\). dword \([0] \leftarrow\) result
VSR[VRT+32]. dwor d [1] \(\leftarrow\) Ox0000_0000_0000_0000
FPSCR. FPRF \(\leftarrow\) ObUUUU
end
FPSCR.FR \(\leftarrow 0\)

Let \(\operatorname{src}\) be the quad-precision floating-point value in VSR[VRB+32].

If \(\operatorname{src}\) is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN and VXCVI are set to 1.

If src is a Quiet NaN or an Infinity, an Invalid Operation exception occurs and VXCVI is set to 1 .

If \(5 r \mathrm{C}\) is a NaN , the result is \(0 \times\) FFFF \(\mathrm{FFFF}_{-} 8000 \_0000\).
Otherwise, if src is a Zero, the result is \(0 \times 0000 \_0000 \_0000 \_0000\).

Otherwise, if \(\operatorname{src}\) is a tInfinity, the result is \(0 \times 0000\) _0000_7FFFFFFF.

Otherwise, if \(\operatorname{src}\) is a - Infinity, the result is OXFFFF_FFFF 8000 _ 0000 .

Otherwise, do the following.
Let rind be the value src truncated to a floating-point integer.

If \(r n d\) is greater than \(+2^{31}\). 1 , an Invalid Operation exception occurs, VXCVI is set to 1 , and the result is \(0 \times 0000\) _0000_7FFF_FFFF.

Otherwise, if r nd is less than \(\cdot 2^{31}\), an Invalid Operation exception occurs, VXCVI is set to 1 , and the result is \(0 \times\) FFFF FFFF \(8000 \_0000\).

Otherwise, the result is the value rnd, and an Inexact exception occurs if \(r\) nd is inexact (i.e., \(r\) nd is not equal to src ).

The result is placed into doubleword element 0 of VSR[ VRT +32 ] in signed integer format.

The contents of doubleword element 1 of VSR[VRT+32] are set to 0 .

FPRF is set to undefined. FR is set to 0 . Fl is set to indicate if the rounded result is inexact.

If an Invalid Operation exception occurs, FR and FI are set to 0 .

If a trap-enabled Invalid Operation exception occurs, VSR[ VRT +32 ] and FPRF are not modified.

See Table 71, "Actions for xscvqpswz," on page 553.

\section*{Special Registers Altered:}

FPRF (undefined) FR (set to 0) Fl
FX VXSNAN VXCVI XX

\section*{VSR Data Layout for xscvqpswz}

VSR[VRB+32]
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|c|}{S1C} \\
\hline VSR[ VRT+32] & \\
\hline tgt.dword[0] & 0x0000_0000_0000_0000 \\
\hline
\end{tabular}


I Table 71. Actions for xscvqpswz

\section*{VSX Scalar truncate \& Convert Quad-Precision format to Unsigned Doubleword format X-form}
xscyqpudz
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 63 & VRT, VRB \\
\hline 0 & 6 & 17 & VRB & & 836 & \((11\) \\
31 \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablel)
reset_xflags()
sic \& bf p_CONERT_FROM_BFP128(VSR[VRB+32])
if src.class. QNaN | src.class. SNaN then do
result \(\leftarrow\) 0x0000_0000_0000_0000
vxsnan_flag \& sic.class. SN̈al
vxcvi_ \(\dagger\) lag \(\leftarrow 1\)
end
else if src.class. Infinity then do
vxcui_flag \(\leftarrow 1\)
if sicisign \(=0\) then
result \(\leftarrow\) OxFFFF_FFFF_FFFF_FFF
else
result \(\leftarrow 0 \times 0000\) _0000 000000000
end
else if src.class. Zero then result \(\leftarrow 0 \times 000000000000 \_0000\)
else do
rnd \(\leftarrow\) bf p_ROUND_TO_ I NTEGER (Ob001, sic)
if bfo_CoMpare_Gi(fnd, \(+2^{64}\). 1 ) then do result \(\leftarrow\) OxFFFF_FFF_FFF_FFFF vxcvi_flag \(\leftarrow 1\)
end
else if bfoc COMPARE LT(rnd, O) then do result \(\leftarrow 0 \times 0000\)-0000_0000_0000 vxcvi_flag \(\leftarrow 1\)
end
else do result \(\leftarrow\) bf p_CoNvert_TO_U164(rnd) if(xx_flag) then Set FX(FPSCR. XX)
    end
end
if(uxsnan_flag) then Setfx(fPSCR.vxSNaN)

vx_flag ↔ vxsnan_flag | vxcvi_flag
ex_flag \(\leftarrow F P S C R, V E \& V x_{-} f l a g\)
if ex_flag=0 then do
VSR[ VRT +32 ]. dword [0] \& result
VSR[VRT+32]. dword [1] \(\leftarrow 0 \times 0000 \_0000 \_0000 \_0000\)
FPSCR. FPRF \(\leftarrow\) ObUUUUU
end
FPSCR. \(F R \leftarrow\left(v x \_f \mid a g=0\right) \& i n c \_f l a g\)
FPSCR.FI \(\leftarrow\left(v x_{-} f \mid a g=0\right) \& x x_{-} f 1 a g\)
Let \(\operatorname{sic}\) be the quad-precision floating-point value in VSR[ VRB +32 ].

If \(\operatorname{src}\) is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN and VXCVI are set to 1.

If src is a Quiet NaN or an Infinity, an Invalid Operation exception occurs and VXCVI is set to 1 .

If \(5 r \mathrm{c}\) is a NaN , the result is \(0 \times 00000_{\mathrm{Z}} 00000_{\mathrm{Z}} 0000 \_0000\).
Otherwise, if src is a Zero, the result is \(0 \times 0000 \_0000 \_0000 \_0000\).

Otherwise, if src is a positive Infinity, the result is OXFFFF_FFF_FFF_FFFF.

Otherwise, if \(\operatorname{src}\) is a negative Infinity, the result is \(0 \times 0000 \_0000 \_0000 \_0000\).

Otherwise, do the following.
Let rnd be the value src truncated to a floating-point integer.

If \(r n d\) is greater than \(+2^{64}\). 1 , an Invalid Operation exception occurs, VXCVI is set to 1 , and the result is \(0 \times F F F F_{\_} F F F F_{-} F F F F_{-} F F F F\).

Otherwise, if rnd is less than 0 , an Invalid Operation exception occurs, VXCVI is set to 1 , and the result is \(0 \times 0000,0000 \_0000 \_0000\).

Otherwise, the result is the value rnd, and an Inexact exception occurs if \(r\) nd is inexact (i.e., \(r\) nd is not equal to src ).

The result is placed into doubleword element 0 of VSR[ VRT +32] in unsigned integer format.

The contents of doubleword element 1 of VSR[VRT+32] are set to 0 .

FPRF is set to undefined. FR is set to 0 . Fl is set to indicate if the rounded result is inexact.

If an Invalid Operation exception occurs, FR and FI are set to 0 .

If a trap-enabled Invalid Operation exception occurs, VSR[ VRT +32 ] and FPRF are not modified.

See Table 72, "Actions for xscvqpudz," on page 555.

\section*{Special Registers Altered:}

FPRF (undefined) FR (set to 0) FI
FX VXSNAN VXCVI XX

\section*{VSR Data Layout for xscvqpudz}

VSR[VRB+32]
\begin{tabular}{l}
\hline \multicolumn{2}{|c|}{ src } \\
\begin{tabular}{|c|c|}
\hline VSR[VRT+32] \\
\hline tgt.dword[0] & \(0 \times 0000 \_0000 \_0000 \_0000\) \\
\hline
\end{tabular}
\end{tabular}

| Table 72. Actions for xscvqpudz

\section*{VSX Scalar truncate \& Convert Quad-Precision format to Unsigned Word format \(X\)-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{xscvqpuwz VRT,VRB} \\
\hline \[
0
\] & \[
\left.\right|_{6} \text { VRT }
\] & 11 & 1 & \[
{ }_{16} \text { VRB }
\] & 21 & 836 & I 1 \\
\hline
\end{tabular}
if MSR. VSX=0 then VSX_Unavailablell
reset_xflags()
sic \&bf p_CONERT_ _ROM_BFP128 (VSR[VRB+32])
if src.class. QNaN | src.class. SNaN then do
result \(\leftarrow\) Ox0000 0000
vxsnan_flag \& sic.class. SNaN
vxcvi_flag \(\leftarrow 1\)
end
else if src.class.Infinity then do
vxcri_flag \(\leftarrow 1\)
if sic.sign \(=0\) then
result \(\leftarrow 0 \times 0000\) _0000_FFFF_FFF
else
result \(\leftarrow 0 \times 0000\) _0000 000000000
end
else if src.class. Zero then
result \(\leftarrow 0 \times 0000\) _0000
else do
ind \& bf p_ROUND_to_ I NTEGER(ObOO1,sic)
if bf cocompare GT(rnd, \(+2^{32} \cdot 1\) ) then do result \(\leftarrow 0 x 0000\) _0000_fFF_ FFFF vxcvi_flag \(\leftarrow 1\)
end
else if bfp_compare_lif(rnd, bfo_zero) then do result \(\leftarrow\) Ox0000_0000_0000_0000 vxcvi_flag \(\leftarrow 1\)
end
else do result \(\leftarrow\) bf p_Convert_to_ul 64(rnd) if(xx_flag) then SetFX(FPSCR. XX)
end
end
if(uxsnan_flag) then Setfx(fpscr.vxsnaw)
if(uxcvi_flag) then SetFX(FPSCR. VXCVI)
vx_flag ャvxsnan_flag | vxcvi_flag
ex_flag \(\leftarrow\) FPSCR. VE \& Ux_flag
if ex_flag=0 then do
VSŔ[VRT+32].dword[0] \& result
VSR[VRT+32]. dwor d [1] \(\leftarrow\) Ox0000_0000_0000_0000
FPSCR. FPRF \(\leftarrow\) ObUUUU
end
FPSCR. \(F R \leftarrow\left(v x \_f \mid a g=0\right) \& i n c \_f\) fag
FPSCR. FI \(\leftarrow\left(v x_{-}^{-} f \mid a g=0\right) \& x x_{-}^{-} \mid a g\)
Let \(\operatorname{src}\) be the quad-precision floating-point value in VSR[VRB+32].

If src is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN and VXCVI are set to 1.

If src is a Quiet NaN or an Infinity, an Invalid Operation exception occurs and VXCVI is set to 1 .

If \(5 r c\) is a NaN , the result is \(0 \times 00000_{-} 0000_{\mathrm{Z}} 0000 \_0000\).
Otherwise, if src is a Zero, the result is \(0 \times 0000 \_0000 \_0000 \_0000\).

Otherwise, if \(\operatorname{src}\) is a positive Infinity, the result is \(0 \times 0000 \_0000\) _FFFFFFFF.

Otherwise, do the following.
Let rnd be the value src truncated to a floating-point integer.

If \(r n d\) is greater than \(+2^{32}\). 1 , an Invalid Operation exception occurs, VXCVI is set to 1 , and the result is \(0 \times 0000\) _ 0000 _FFFF_FFFF.

Otherwise, if rnd is less than 0 , an Invalid Operation exception occurs, VXCVI is set to 1 , and the result is \(0 \times 0000 \_0000 \_0000 \_0000\).

Otherwise, the result is the value rnd, and an Inexact exception occurs if \(r\) nd is inexact (i.e., \(r\) nd is not equal to src ).

The result is placed into doubleword element 0 of VSR[ VRT+32] in unsigned integer format.

The contents of doubleword element 1 of VSR[VRT+32] are set to 0 .

FPRF is set to undefined. FR is set to 0 . Fl is set to indicate if the rounded result is inexact.

If an Invalid Operation exception occurs, \(F R\) and \(F \mid\) are set to 0 .

If a trap-enabled Invalid Operation exception occurs, VSR[ VRT +32 ] and FPRF are not modified.

See Table 73, "Actions for xscvqpuwz," on page 557.

\section*{Special Registers Altered:}
```

        FPRF (undefined) FR (set to 0) FI
    ```
    FX VXSNAN VXCVI XX
| VSR Data Layout for xscvqpuwz
I VSR[VRB+32]
I
I
VSR[ VRT+32]
\(\square\)


I Table 73. Actions for xscvqpuwz

\section*{Version 3.0}

VSX Scalar Convert Signed Doubleword
format to Quad-Precision format X-form
xscvsdqp VRT,VRB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \[
63
\] & \[
{ }_{6} \mathrm{VRT}
\] & \[
10
\] & \[
{ }_{16} \text { VRB }
\] & 21 & 836 & \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablell
sTc \(\leftarrow\) bfp_CONVERT_FROM_SI 64(VSR[VRB+32], dword[0])
result \(\leftarrow\) bfp_CONVERT_TO_BFP128(src)
VSR[VRT+32] \(\leftarrow\) result
FPSCR.FPRF \(\leftarrow\) fprf_CLASS_BFP128(result)
FPSCR.FR \(\leftarrow 0\)
FPSCR.FI \(\leftarrow 0\)
Let src be the signed integer value in doubleword element 0 of VSR[ VRB +32 ].

SrC is placed into VSR[VRT+32] in quad-precision floating-point format.

FPRF is set to the class and sign of the result. \(F R\) is set to \(0 . \mathrm{Fl}\) is set to 0 .

Special Registers Altered: FPRF FR (set to 0) FI (set to 0)
| VSR Data Layout for xscvsdqp
VSR[ VRB+32]
\begin{tabular}{|c|c|}
\hline src.dword[0] & unused \\
\hline
\end{tabular}
I VSR[VRT+32]
I \(\square\)

\section*{VSX Scalar Convert Single-Precision to Double-Precision format XX2-form}
xscvspdp XT,XB
\begin{tabular}{|ll|l|l|l|l|l|l|l|}
\hline 60 & & T & & III & & B & & 329 \\
\hline 0 & & 6 & & 11 & & 16 & & 21 \\
\hline
\end{tabular}
```

```
    reset_xflags()
```

```
    reset_xflags()
    src \leftarrowVSR[32xBX+B].word[0]
    src \leftarrowVSR[32xBX+B].word[0]
    result }\leftarrow\mathrm{ ConvertVectorSPtoScalarSP(src)
    result }\leftarrow\mathrm{ ConvertVectorSPtoScalarSP(src)
    if(vxsnan_flag) then SetFX(FPSCR.VXSNAN)
    if(vxsnan_flag) then SetFX(FPSCR.VXSNAN)
    vex_flag \leftarrow FPSCR.VE & vxsnan_flag
    vex_flag \leftarrow FPSCR.VE & vxsnan_flag
    FPSCR.FR \leftarrow0b0
    FPSCR.FR \leftarrow0b0
    FPSCR.FI \leftarrow0b0
    FPSCR.FI \leftarrow0b0
    if( ~vex_flag) then do
    if( ~vex_flag) then do
        VSR[32xTX+T].dword[0] }\leftarrow\mathrm{ result
        VSR[32xTX+T].dword[0] }\leftarrow\mathrm{ result
        VSR[32xTX+T].dword[1] & 0xUUUU_UUUU_UUUU_UUUU
        VSR[32xTX+T].dword[1] & 0xUUUU_UUUU_UUUU_UUUU
        FPSCR.FPRF \leftarrow ClassDP(result)
        FPSCR.FPRF \leftarrow ClassDP(result)
    end
```

```
    end
```

```
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).

Let src be the single-precision floating-point value in word element 0 of VSR[ XB].

If src is a SNaN , the result is src, converted to a QNaN (i.e., bit 9 of src set to 1 ). VXSNAN is set to 1 .

Otherwise, the result is src .

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. \(F R\) is set to \(0 . \mathrm{Fl}\) is set to 0 .

If a trap-enabled invalid operation exception occurs, VSR[XT] is not modified, FPRF is not modified, \(F R\) is set to 0 , and Fl is set to 0 .

\section*{Special Registers Altered}
```

FPRF FR=0bO Fl=0bO FX
FX VXSNAN

```
\(\operatorname{src}=\) VSR[ XB]
QNaN (i.e., bit 9 of \(\operatorname{src}\) set to 1 ). VXSNAN is set to 1 .

\section*{VSR Data Layout for xscvspdp \\ VSR Data Layout for xscvspdp}
src = VSR[ XB]
\begin{tabular}{|c|c|c|}
\hline .word[0] & unused & unused \\
\hline
\end{tabular}
tgt = VSR[XT]
\begin{tabular}{|c|l|}
\hline &.\(d\) word[0] \\
\hline 0 & undefined \\
\hline
\end{tabular}

\section*{Programming Note}
xscuspdp can be used to convert a single-precision value in single-precision format to double-precision format for use by Floating-Point scalar single-precision operations.

\section*{Version 3.0}

VSX Scalar Convert Single-Precision to Double-Precision format Non-signalling XX2-form
xscvspdpn \(\quad \mathrm{XT}, \mathrm{XB}\)
\begin{tabular}{|c|cc|c|c|c|c|}
\hline 60 & & T & & III & & B \\
0 & & 6 & & 11 & & \\
\hline
\end{tabular}
reset_xflags()
src \(\leftarrow \operatorname{VSR}[32 \times B X+B]\). word[0]
result \(\leftarrow\) ConvertSPtodP_NS(src)
\(\operatorname{VSR}[32 \times T X+T]\). dword \([0] \leftarrow\) result
VSR[32xTX+T].dword[1] \& 0xUUUU_UUUU_UUUU_UUUU
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).

Let src be the single-precision floating-point value in word element 0 of VSR[ XB].
\(\operatorname{src}\) is placed into doubleword element 0 of VSR[ XT] in double-precision format.

The contents of doubleword element 1 of VSR[ XT] are undefined.

\section*{Special Registers Altered \\ None}

\section*{VSR Data Layout for xscvspdpn}
sic = VSR[ XB]
\begin{tabular}{|c|c|c|c|}
\hline .word[0] & unused & unused & unused \\
\hline
\end{tabular}
\(\operatorname{tgt}=\operatorname{VSR}[X T]\)
\begin{tabular}{|c|cc|}
\hline & dword[0] & undefined \\
\hline 0 & 32 & 64
\end{tabular}

\section*{Programming Note}
xscvspdp should be used to convert a vector single-precision floating-point value to scalar double-precision format.
xscvspdpn should be used to convert a vector single-precision floating-point value to scalar single-precision format.

\section*{VSX Scalar Convert Signed Integer Doubleword to floating-point format and round to Double-Precision format XX2-form}
\[
\text { xscvsxddp } \quad \text { XT,XB }
\]
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline -60 & 6 & T & 11 & III & 16 & B & 21 & 376 & Bx 8 TY \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline XT & \(\leftarrow \mathrm{TX} \|\) \\
\hline XB & \(\leftarrow B X \|\) \\
\hline & \\
\hline
\end{tabular}
reset_xflags()
v\{0:inf\} \(\leftarrow\) ConvertSDtoFP(VSR[XB]\{0:63\})
result \(\{0: 63\} \leftarrow \operatorname{RoundToDP}(\mathrm{RN}, \mathrm{v})\)
VSR[XT] \(\leftarrow\) result || 0xUUUU_UUUU_UUUU_UUUU
if(xx_flag) then \(\operatorname{SetFX}(X X)\)
FPRF \(\quad \leftarrow\) ClassDP(result)
FR \(\quad \leftarrow\) inc_flag
FI \(\quad \leftarrow x x_{-} f l a g\)
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the signed integer value in doubleword element 0 of \(\operatorname{VSR}[X B]\).
src is converted to an unbounded-precision floating-point value and rounded to double-precision using the rounding mode specified by RN.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. FI is set to indicate the result is inexact.


\section*{VSX Scalar Convert Signed Integer Doubleword to floating-point format and round to Single-Precision XX2-form}
xscvsxdsp XT,XB

```

reset_xflags()
src \leftarrow ConvertSDtoDP(VSR[32xBX+B].dword[0])
result }\leftarrow\mathrm{ RoundToSP(RN,src)
VSR[32\timesTX+T].dword[0] \leftarrow ConvertSPtoSP64 (result)
VSR[32xTX+T].dword[1] \& OxUUUU_UOUU_UOU__UUU
if(xx_flag) then SetFX(XX)
FPRF }\leftarrow\mathrm{ ClasSSP(result)
FR }\leftarrow\mathrm{ inc_flag
FI }\leftarrow\textrm{xx_flag

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the two's-complement integer value in doubleword element 0 of \(\mathrm{VSR}[\mathrm{XB}]\).
src is converted to floating-point format, and rounded to single-precision using the rounding mode specified by RN.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to indicate if the result was incremented when rounded. FI is set to indicate the result is inexact.
```

Special Registers Altered
FPRF FR FI FX XX

```

VSR Data Layout for xscvsxdsp src = VSR[XB]
\begin{tabular}{l}
\hline \multicolumn{2}{|c|}{ SD } & unused \\
\hline tgt = VSR[XT] \\
\hline DP \\
\hline 0
\end{tabular}

VSX Scalar Convert Signed Doubleword format to Quad-Precision format X-form
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{xsclvsdqp VRT,VRB} \\
\hline \[
\begin{array}{ll} 
& 63 \\
0
\end{array}
\] & \[
{ }_{6} \mathrm{VRT}
\] & \[
10
\] & \[
{ }_{16} \text { VRB }
\] & \[
\begin{array}{ll}
\hline & 836 \\
21 &
\end{array}
\] & \begin{tabular}{|r|}
1 \\
31
\end{tabular} \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablell
sic \(\leftarrow\) bfp_CONVERT_FROM_SI64(VSR[VRB+32], dword[0])
result \(\leftarrow\) bfp_CONVERT_TO_BFP128(src)
VSR[VRT+32] \(\leftarrow\) result
FPSCR.FPRF \(\leftarrow\) fprf_CLASS_BFP128(result)
FPSCR.FR \(\leftarrow 0\)
FPSCR.FI \(\leftarrow 0\)
Let \(\operatorname{src}\) be the signed integer value in doubleword element 0 of VSR[VRB+32].
\(\operatorname{src}\) is placed into VSR[VRT+32] in quad-precision floating-point format.

FPRF is set to the class and sign of the result. \(F R\) is set to \(0 . \mathrm{Fl}\) is set to 0 .

Special Registers Altered:
FPRF FR (set to 0) FI (set to 0)
| VSR Data Layout for xscvsdqp
- VSR[VRB+32]
\begin{tabular}{|l|c|}
\hline src. dword[0] & unused \\
\hline I VSR[VRT +32\(]\) \\
\hline tgt \\
\hline
\end{tabular}

VSX Scalar Convert Unsigned Doubleword format to Quad-Precision format X-form
xscyudqp
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 63 & VRT \\
\hline 6 & & 6 & & 2 & VRB & \\
\hline 11 & & 16 & 21 & & & \\
31 \\
\hline
\end{tabular}
```

If MSR.VSX=0 then VSX_Unavailablell
src \&bfp_CONVERT_FROM_Ul 64(VSR[VRB+32],dword[0])
result \&bfp_CONVERT_TO_BFP128(src)
VSR[VRT+32]}\leftarrow\mathrm{ result
FPSCR,FPRF \&fprf_CLASS_BFP128(result)
FPSCR.FR }\leftarrow
FPSCR.FI }\leftarrow

```

Let \(\operatorname{sic}\) be the unsigned integer value in doubleword element 0 of VSR[ VRB +32 ].

SrC is placed into VSR[VRT+32] in quad-precision floating-point format.

FPRF is set to the class and sign of the result. FR is set to \(0 . \mathrm{Fl}\) is set to 0 .

Special Registers Altered:
FPRF FR (set to 0) FI (set to 0)
| VSR Data Layout for xscvudqp
I VSR[VRb+32]
\(\square\)
I \begin{tabular}{|c|c|}
\hline src.dword[0] & unused \\
\hline
\end{tabular}
I VSR[VRT+32]
I \(\square\)

\section*{VSX Scalar Convert Unsigned Integer Doubleword to floating-point format and round to Double-Precision format XX2-form}
\[
\text { xscvuxddp } \quad \mathrm{XT}, \mathrm{XB}
\]
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 60 & & T & & III & & B & \\
\hline
\end{tabular}
\begin{tabular}{ll}
\(X T\) & \(\leftarrow T X \| T\) \\
\(X B\) & \(\leftarrow B X \| B\)
\end{tabular}
reset_xflags()
src\{0:inf\} \(\leftarrow\) ConvertUDtoFP(VSR[XB]\{0:63\})
result \(\{0: 63\} \leftarrow\) RoundToDP(RN, src)
VSR[XT] \(\leftarrow\) result || 0xUUUU_UUUU_UUUU_UUUU
if(xx_flag) then \(\operatorname{SetFX}(X X)\)
FPRF \(\quad \leftarrow\) ClassDP(result)
FR \(\quad \leftarrow\) inc_flag
FI \(\quad \leftarrow x x_{-} f l a g\)
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the unsigned integer value in doubleword element 0 of VSR[XB].
src is converted to an unbounded-precision floating-point value and rounded to double-precision using the rounding mode specified by RN.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. FI is set to indicate the result is inexact.

Special Registers Altered
```

    FPRF FR FI FX XX
    ```

VSR Data Layout for xscvuxddp
src = VSR[XB]
\begin{tabular}{|l|l|}
\hline \multicolumn{2}{|c|}{ UD } \\
\hline tgt \(=\) VSR[XT] & unused \\
\hline \begin{tabular}{|l|l|}
\hline DP & undefined \\
\hline 0 & 64
\end{tabular} & 127 \\
\hline
\end{tabular}

\section*{VSX Scalar Convert Unsigned Integer Doubleword to floating-point format and round to Single-Precision XX2-form}
\[
\text { xscvuxdsp } \quad \text { XT,XB }
\]
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 60 & & T & & III & & B \\
\hline 0 & & 6 & & 11 & & 296 \\
\hline 16 & & 21 & & \\
\hline
\end{tabular}
```

reset_xflags()
src \leftarrow ConvertUDtoDP(VSR[32xBX+B].dword[0])
result }\leftarrow\mathrm{ RoundToSP(RN,src)
VSR[32\timesTX+T].dword[0] \leftarrow ConvertSPtoSP64 (result)
VSR[32xTX+T] .dword[1] \& OxUUOU_UUUU_UUUU_UOU
if(xx_flag) then SetFX(XX)
FPRF }\leftarrow\mathrm{ ClasSSP(result)
FR}\leftarrow\mathrm{ inc_flag
FI }\leftarrow\textrm{xx_flag

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the unsigned-integer value in doubleword element 0 of VSR[XB].
src is converted to floating-point format, and rounded to single-precision using the rounding mode specified by RN.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.
```

Special Registers Altered
FPRF FR FI FX XX

```

VSR Data Layout for xscvuxdsp src = VSR[XB]
\begin{tabular}{l}
\hline \multicolumn{2}{|c|}{ UD } & unused \\
\hline tgt = VSR[XT] \\
\hline DP \\
\hline 0
\end{tabular}

VSX Scalar Divide Double-Precision XX3-form
xsdivdp XT,XA,XB
\begin{tabular}{|c|cc|c|c|c|c|c|c|}
\hline 60 & & T & & A & & B & & 56 \\
\hline 0 & & 6 & & 11 & & 16 & & 21
\end{tabular}
\begin{tabular}{|c|c|}
\hline XT & \(\leftarrow \mathrm{TX} \| \mathrm{T}\) \\
\hline XA & \(\leftarrow A X \| A\) \\
\hline ХВ & \(\leftarrow B X|\mid B\) \\
\hline \multicolumn{2}{|l|}{reset_xflags()} \\
\hline & \(\leftarrow \mathrm{VSR}[\mathrm{XA}]\{0: 63\}\) \\
\hline src2 & \(\leftarrow \mathrm{VSR}[\mathrm{XB}]\{0: 63\}\) \\
\hline v\{0:inf \(\}\) & \(\leftarrow\) DivideFP(src1, src2) \\
\hline \multicolumn{2}{|l|}{result \(\{0: 63\} \leftarrow \operatorname{RoundToDP}(\mathrm{RN}, \mathrm{v})\)} \\
\hline if(vxsnan_flag) & g) then SetFX(VXSNAN) \\
\hline if(vxidi_flag) & ) then SetFX(VXIDI) \\
\hline if(vxzdz_flag) & ) then SetFX(VXZDZ) \\
\hline if(ox_flag) & then SetFX(0X) \\
\hline if(ux_flag) & then SetFX(UX) \\
\hline if(xx_flag) & then SetFX(XX) \\
\hline if(zx_flag) & then \(\operatorname{SetFX}(\mathrm{ZX})\) \\
\hline vex_flag & \(\leftarrow\) VE \& (vxsnan_flag | vxidi_flag | vxzdz_flag) \\
\hline zex_flag & \(\leftarrow\) ZE \& zx_flag \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{if( ~vex_flag \& ~zex_flag ) then do} \\
\hline VSR[XT] & = result || 0xUUUU_UUUU_UUUU_UUUU \\
\hline FPRF & = ClassDP(result) \\
\hline FR & = inc_flag \\
\hline FI & = xx_flag \\
\hline \multicolumn{2}{|l|}{end} \\
\hline \multicolumn{2}{|l|}{else do} \\
\hline FR & \(=0 \mathrm{bo}\) \\
\hline FI & = 0b0 \\
\hline end & \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

Let src2 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[\mathrm{XB}]\).
src1 is divided \({ }^{[1]}\) by src2, producing a quotient having unbounded range and precision.

The quotient is normalized \({ }^{[2]}\).
See Actions for xsdivdp (p. 565).
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception or a trap-enabled zero divide exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0.

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered}

FPRF FR FI FX OX UX ZX XX
VXSNAN VXIDI VXZDZ
VSR Data Layout for xsdivdp
src1 = VSR[XA]
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline src2 \(=\mathrm{VSR}[\mathrm{XB}]\) & \\
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|ll|}
\hline DP & \multicolumn{2}{c|}{ undefined } \\
\hline 0 & 64 & 127 \\
\hline
\end{tabular}

\footnotetext{
1. Floating-point division is based on exponent subtraction and division of the significands.
2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \[
\begin{aligned}
& v \leftarrow d Q N a N \\
& \text { vxidi_flag } \leftarrow 1
\end{aligned}
\] & \(v \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow-\) Infinity & \(V \leftarrow\) - Infinity & \[
\begin{array}{|l}
\hline v \leftarrow d Q N a N \\
\text { vxidi_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline -NZF & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{V} \leftarrow+\ln f i n i t y\) Zx_flag \(\leftarrow 1\) & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow-\operatorname{lnfinity} \\
& \mathrm{zx} \text { flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{V} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\
\mathrm{vxsnan} \text { _flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline -Zero & \(\mathrm{V} \leftarrow+\) Zero & \(v \leftarrow+\) Zero & \[
\begin{aligned}
& v \leftarrow d Q N a N \\
& \text { vxzdz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \mathrm{vxzdz} \text { _lag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{V} \leftarrow\)-Zero & \(v \leftarrow\)-Zero & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\overline{\text { - }}\) +Zero & \(\mathrm{V} \leftarrow\)-Zero & \(v \leftarrow\)-Zero & \[
\begin{array}{|l|}
\hline v \leftarrow d Q N a N \\
\text { vxzdz_lag } \leftarrow 1 \\
\hline
\end{array}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \mathrm{vxzdz} \text { _lag } \leftarrow 1
\end{aligned}
\] & \(v \leftarrow+\) Zero & \(v \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l}
\hline v \leftarrow Q(\text { src2 }) \\
v x \text { nnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline ¢ +NZF & \(\mathrm{V} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow-\text { Infinity } \\
& \text { zx_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(V \leftarrow+\) Infinity zx_flag \(\leftarrow 1\) & \(\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)\) & \(v \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline +Infinity & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{dQNaN} \\
& \text { vxidi_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{V} \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow \mathrm{dQNaN} \\
& \text { vxidi_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline QNaN & \(\mathrm{V} \leftarrow \mathrm{src} 1\) & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \[
\begin{aligned}
& v \leftarrow \text { src1 } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline SNaN & \begin{tabular}{l}
\[
\mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1)
\] \\
vxsnan_flag \(\leftarrow 1\)
\end{tabular} & \[
\begin{aligned}
& v \leftarrow Q \text { (src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& v \leftarrow Q(\text { src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { src1 } 1) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& v \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& v \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& V \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { src1 }) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline Explanation: & & & & & & & & \\
\hline src1 & \multicolumn{8}{|l|}{The double-precision floating-point value in doubleword element 0 of VSR[XA].} \\
\hline src2 & \multicolumn{8}{|l|}{\multirow[t]{2}{*}{The double-precision floating-point value in doubleword element 0 of VSR[XB].}} \\
\hline dQNaN & & & & & & & & \\
\hline NZF & \multicolumn{8}{|l|}{\begin{tabular}{l}
Default quiet NaN (0x7FF8_0000_0000_0000). \\
Nonzero finite number.
\end{tabular}} \\
\hline \(D(x, y)\) & \multicolumn{8}{|l|}{\multirow[t]{2}{*}{Return the normalized quotient of floating-point value x divided by floating-point value y , having unbounded range and precision. Return a QNaN with the payload of \(x\).}} \\
\hline \(Q(x)\) & & & & & & & & \\
\hline \(\checkmark\) & \multicolumn{8}{|l|}{The intermediate result having unbounded signficand precision and unbounded exponent range.} \\
\hline
\end{tabular}

Table 74.Actions for xsdivdp

VSX Scalar Divide Quad-Precision [using round to Odd] X-form
\begin{tabular}{l}
\begin{tabular}{l} 
xsdivqp \\
xsdivqpo
\end{tabular} \\
\begin{tabular}{|c|c|c|c|c|c|} 
VRT,VRA,VRB \\
VRT,VRA,VRB
\end{tabular} \\
\hline 63 \\
\hline 0
\end{tabular}
if MSR. VSX=O then VSX_Unavailablel)
reset_xflags()
sicl \(\leftarrow\) bfo_Convert_from_bpp128(VSR[Vra+32])
sic2 \& bfoconvert_FROM_BFP128(VSR[VBB+32])
\(v \quad \leftarrow\) bfo_olvide(sicl, sicz)
rnd \(\leftarrow\) bfp_ROUND_TO_BFP128(RO, FPSCR.RN, v)
result \(\leftarrow\) bf poconvert_TO_BFP128(rnd)
if(uxsnan_flag) then Set \(\operatorname{PX}(F P S C R\). VXSNAN)
if(uxidi_flag) then SetFX(FPSCR. VXIDI)
if(vxzdz_flag) then SetFX(FPSCR. VXZDZ)
if(ox_flag) then SetFX(FPSCR. OX)
if(ux_flag) then SetfX(FPSCR. UX)
if(2x_flag) then SetfX(FPSCR. ZX)
if(xx_- \(f\) ag \()\) then SetFX(FPSCR. XX)
vx_flag ↔vxsnan_flag|vxidi_flag|vxzdz_flag

if ex fiag=0 then do
VsR[VRT +32\(] \leftarrow\) result
FPSCR. FPRF \& forf_Class_bFP128(result)
end

FPSCR. \(F 1 \leftarrow\left(v x_{-}^{-} f \mid a g=0\right) \&\left(2 x_{-}^{-} f \mid a g=0\right) \& x x_{-} \hat{f} \mid a g\)
Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let src2 be the floating-point value in VSR[VRB+32] represented in quad-precision format.

If either srcl or srcl is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN is set to 1

If srcl and srct are Infinity values, an Invalid Operation exception occurs and VXIDI is set to 1 .

If \(\mathrm{srcl}_{\mathrm{Cl}}\) and \(\mathrm{srcc}^{2}\) are Zero values, an Invalid Operation exception occurs and VXZDZ is set to 1 .

If \(\operatorname{srcl}\) is a finite value and \(\operatorname{srch}^{2}\) is a Zero value, an Zero Divide exception occurs and \(Z X\) is set to 1 .

If \(\operatorname{srcl}\) is a Signalling NaN , the result is the Quiet NaN corresponding to srcl.

Otherwise, if srcl is a Quiet NaN , the result is srcl .
Otherwise, if \(\mathrm{src}_{\mathrm{C}}\) is a Signalling NaN, the result is the Quiet NaN corresponding to sic2.

Otherwise, if \(s r^{\prime} 2\) is a Quiet NaN , the result is src .
Otherwise, if \(\mathrm{srCl}^{2}\) and \(\mathrm{srcl}_{2}\) are Infinity values, or if srcl and \(\mathrm{srcl}_{2}\) are Zero values, the result is the default Quiet \(\mathrm{NaN}^{[1]}\).

Otherwise, if srcl is a non-zero value and \(\operatorname{srch}\) is a Zero value, the result is an Infinity.

Otherwise, do the following.
The normalized quotient of srcl divided by srcc is produced with unbounded significand precision and exponent range.

See Table 75, "Actions for xsdivqp[o]," on page 567.

If the intermediate result is Tiny (i.e., the unbiased exponent is less than \(\cdot 16382\) ) and \(U E=0\), the significand is shifted right \(N\) bits, where \(N\) is the difference between - 16382 and the unbiased exponent of the intermediate result. The exponent of the intermediate result is set to the value - 16382.

If \(R O=1\), let the rounding mode be Round to Odd. Otherwise, let the rounding mode be specified by RN. Unless the result is an Infinity or a Zero, the intermediate result is rounded to quad-precision using the specified rounding mode.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into VSR[ VRT +32 ] in quad-precision format.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-disabled Invalid Operation exception occurs, \(F R\) and \(F I\) are set to 0 .

If a trap-disabled Zero Divide exception occurs, FR and Fl are set to 0 .

If a trap-enabled Invalid Operation exception or a trap-enabled Zero Divide exception occurs, VSR[VRT+32] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered:}

FPRF FR FI
FX VXSNAN VXIDI VXZDZ OX UX ZX XX

\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{9}{|l|}{VSR Data Layout for xsdivqp[o]} \\
\hline \multicolumn{9}{|l|}{VSR[ VRA +32 ]} \\
\hline \multicolumn{4}{|c|}{srcl} & & & & & \\
\hline \multicolumn{4}{|l|}{VSR[ VRB +32 ]} & & & & & \\
\hline \multicolumn{4}{|c|}{sic2} & & & & & \\
\hline \multicolumn{4}{|l|}{VSR[ VRT+32]} & & & & & \\
\hline \multicolumn{4}{|c|}{tgt} & & & & & \\
\hline \multicolumn{9}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity &  & \(v \leftarrow\) Hntinity & \(v \leftarrow H\) finity & \(v \leftarrow \cdot \mid\) nfinity & \(v \leftarrow\) Infinity &  & \multirow{6}{*}{\(1+\) sicl} & \multirow{6}{*}{\(v \leftarrow\) quiet \(\mid\) sicz wxsmon_flag \(\leftarrow 1\)} \\
\hline -NZF & \multirow[b]{2}{*}{\(v \leqslant\) Hero} & \(v\) - Div(sicl, scicl & \[
\begin{gathered}
1 \leftarrow+1 n \text { finity } \\
2 x . f 12 g+1 \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
v \leftarrow \cdot \operatorname{lnininity} \\
x \times \| \operatorname{lig} \leftarrow 1
\end{gathered}
\] & \(v \leftarrow \operatorname{Div}(\mathrm{scc}, \mathrm{sc}(2)\) & \multirow[b]{2}{*}{\(1 \leqslant\) lero} & & \\
\hline -Zero & & & \multicolumn{2}{|c|}{\multirow[b]{2}{*}{}} & & & & \\
\hline - +Zero & \multirow[t]{2}{*}{\(1 \leqslant\) lero} & & & & & \multirow[t]{2}{*}{\(v \leqslant\) lero} & & \\
\hline +NZF & & \(v \leftarrow \operatorname{Div}(\operatorname{sic}\), sic 2 ) & \[
\begin{gathered}
v \leftarrow \cdot \operatorname{lnfinity} \\
2 x . f \log +1 \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
\hline v \leftarrow+\mid n f i n i t y \\
2 x .1 \operatorname{lig}+1
\end{gathered}
\] & \(v \leftarrow \operatorname{Div}(\mathrm{sic}\), sict 2\()\) & & & \\
\hline +Infinity & \[
\begin{gathered}
V+\text { dollill } \\
\text { vxidif } \mathrm{H} \text { lgg }+1 \\
\hline
\end{gathered}
\] & \(v \leftarrow \cdot \mid n f i n i t y\) & \(v \leftarrow \cdot \mid n t i n i t y\) & \(v \leftarrow+1\) finity & \(v \leftarrow\) Hafinity & \[
\begin{gathered}
v \leftarrow \text { dillal } \\
\text { wxidif flag } \leftarrow 1 \\
\hline
\end{gathered}
\] & & \\
\hline QNaN & \multicolumn{4}{|r|}{1 +sicl} & & & & \[
\begin{gathered}
V \leftarrow \text { sicl } \\
\text { vxsnan flag } \leftarrow 1
\end{gathered}
\] \\
\hline SNaN & \multicolumn{8}{|c|}{\[
\begin{aligned}
& V \leftarrow \text { quiet }(\text { scri] } \\
& \text { vxssan_f } 1 \text { ag }+1
\end{aligned}
\]} \\
\hline \multicolumn{9}{|l|}{Explanation:} \\
\hline stcl & \multicolumn{8}{|l|}{he quad-precision floating-point value in VSR[ VRA +32\(]\).} \\
\hline sic2 & \multicolumn{8}{|l|}{The quad-precision floating-point value in VSR[ VRB +32].} \\
\hline dQNaN & \multicolumn{8}{|l|}{\multirow[t]{2}{*}{Default quiet NaN (0x7FFF \(\left.8000 \_0000 \_0000 \_0000 \_0000 \_0000\right)\).}} \\
\hline \multirow[t]{2}{*}{\[
\begin{aligned}
& \text { NZF } \\
& \text { Div( } x, y)
\end{aligned}
\]} & & & & & & & & \\
\hline & \multicolumn{8}{|l|}{The floating-point value x is divided \({ }^{1}\) by floating-point value y . Return the normalized \({ }^{2}\) quotient, having unbounded range and precision.} \\
\hline \multirow[t]{2}{*}{quiet ( x )} & \multicolumn{8}{|l|}{Convert x to the corresponding Quiet NaN.} \\
\hline & \multicolumn{8}{|l|}{The intermediate result having unbounded significand precision and unbounded exponent range.} \\
\hline
\end{tabular}

\section*{I Table 75. Actions for xsdivqp[o]}
1. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then subtracted or added as appropriate, depending on the signs of the operands, to form an intermediate difference. All 64 bits of the significand as well as all three guard bits ( \(G, R\), and \(X\) ) enter into the computation.
2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.

VSX Scalar Divide Single-Precision XX3-form
\[
\text { xsdivsp } \quad X T, X A, X B
\]

```

reset_xflags()
src1 }\leftarrow\operatorname{VSR[32\timesAX+A].dword[0]
Src2 }\leftarrow\textrm{VSR[32\timesBX+B].dword[0]
v }\leftarrow\mathrm{ DivideDP(src1,src2)
result }\leftarrow\mathrm{ RoundToSP(RN,v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxidi_flag) then SetFX(VXIDI)
if(vxzdz_flag) then SetFX(VXZDZ)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
if(zx_flag) then SetFX(ZX)
vex_flag \leftarrowVE \& (vxsnan_flag|vxidi_flag|vxzdz_flag)
zex_flag \leftarrow ZE \& zx_flag
if( ~vex_flag \& ~zex_flag ) then do
VSR[32xTX+T].dword[0] \leftarrow ConvertSPtoSP64 (result)
VSR[32xTX+T].dword[1] \& 0xUUUU_UUUU_UUUU_UUUU
FPRF}\leftarrowClassSP(result
FR }\leftarrow\mathrm{ inc_flag
FI }\leftarrowxx_fla
end
else do
FR}\leftarrow0.\textrm{bO
FI }\leftarrow0\textrm{ObO
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src1 is divided \({ }^{[1]}\) by src2, producing a quotient having unbounded range and precision.

The quotient is normalized \({ }^{[2]}\).
See Table 76, "Actions for xsdivsp," on page 569.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.
If a trap-enabled invalid operation exception or a trap-enabled zero divide exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0.

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

Special Registers Altered
FPRF FR FI FX OX UX ZX XX
VXSNAN VXIDI VXZDZ

\section*{VSR Data Layout for xsdivsp}
src1 = VSR[XA]
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline src2 \(=\) VSR[XB] & \\
\hline DP & unused \\
\hline tgt \(=\) VSR[XT] & undefined \\
\hline DP & \\
\hline 0 & 64
\end{tabular}

\footnotetext{
1. Floating-point division is based on exponent subtraction and division of the significands.
2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vxidi__lag } \leftarrow 1
\end{aligned}
\] & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow-\ln\) finity & \(\mathrm{V} \leftarrow-\) Infinity & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vxidi_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \begin{tabular}{l}
\[
\mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2)
\] \\
vxsnan_flag \(\leftarrow 1\)
\end{tabular} \\
\hline -NZF & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)\) & \begin{tabular}{l}
\(\mathrm{V} \leftarrow+\) Infinity \\
zx_flag \(\leftarrow 1\)
\end{tabular} & \[
\begin{aligned}
& \mathrm{V} \leftarrow-\operatorname{lnfinity} \\
& \text { zx_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{V} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline -Zero & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow+\) Zero & \[
\begin{array}{|l|}
\hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
\text { vxzdz_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \[
\begin{array}{|l|}
\hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
\text { vxzdz_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \(\mathrm{V} \leftarrow-\) Zero & \(\mathrm{v} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline - + Zero & \(\mathrm{V} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow\)-Zero & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{dQNaN} \\
& \text { vxzdz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{array}{|l|}
\hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
\mathrm{vxzdz} \text { _lag } \leftarrow 1 \\
\hline
\end{array}
\] & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{v} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline ¢ +NZF & \(\mathrm{V} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)\) & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow-\operatorname{lnfinity} \\
& \mathrm{zx} \text { _lag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{V} \leftarrow+\) Infinity Zx_flag \(\leftarrow 1\) & \(\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline +Infinity & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{dQNaN} \\
& \text { vxidi_flag } \leftarrow 1
\end{aligned}
\] & \(V \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow-\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{dQNaN} \\
& \text { vxidi_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline QNaN & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{SrCl} 1 \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline SNaN & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src1}) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& v \leftarrow Q(\operatorname{src} 1) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{array}{|l}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& v \leftarrow Q(\operatorname{src} 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{array}{|l}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] \\
\hline \multicolumn{9}{|l|}{Explanation:} \\
\hline src1 & \multicolumn{8}{|l|}{The double-precision floating-point value in doubleword element 0 of VSR[XA].} \\
\hline src2 & \multicolumn{8}{|l|}{The double-precision floating-point value in doubleword element 0 of VSR[XB].} \\
\hline dQNaN & \multicolumn{8}{|l|}{Default quiet NaN (0x7FF8_0000_0000_0000).} \\
\hline NZF & \multicolumn{8}{|l|}{Nonzero finite number.} \\
\hline \(D(x, y)\) & \multicolumn{8}{|l|}{Return the normalized quotient of floating-point value x divided by floating-point value y , having unbounded range and precision.} \\
\hline Q(x) & \multicolumn{8}{|l|}{Return a QNaN with the payload of \(x\).} \\
\hline \(\checkmark\) & \multicolumn{8}{|l|}{The intermediate result having unbounded signficand precision and unbounded exponent range.} \\
\hline
\end{tabular}

Table 76.Actions for xsdivsp

\section*{Version 3.0}

\section*{VSX Scalar Insert Exponent Double-Precision X-form}
xsiexpdp \begin{tabular}{l} 
XT,RA,RB \\
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & RA & RB & & 918 \\
0 & & 6 & & 11 & & \\
\hline 12
\end{tabular}\(|\)
\end{tabular}
if MSR. VSX=0 then VSX_Unavailablell
\(\operatorname{sicl} \leftarrow\) GPR[RA]
\(\operatorname{sic} 2 \leftarrow G P R[R B]\)
VSR[32xTXXT].dword[0].bit[0] \(\leftarrow\) sicl.bit[0]
VSR[32xTXXT], dword[0], bit[1:11] \(\leftarrow\) src2. bit [53:63]
VSR[32xTX+T], dword[0], bit[12:63] \(\operatorname{sic} 1\), bit [12: 63]
VSR[32xTX+T]. dwo od [1] \(\leftarrow\) OxUUU__UUUZ_UUUU_UUUUU

Let \(X T\) be the sum \(32 \times T X+T\).
Let \(\operatorname{srcl}\) be the unsigned integer value in GPR[RA].
Let \(\operatorname{src} 2\) be the unsigned integer value in GPR[RB].
The contents of bit 0 of \(\mathrm{srcl}^{\text {are placed into bit } 0 \text { of }}\) VSR[ XT].

The contents of bits 53:63 of src2 are placed into bits 1:11 of VSR[ XT].

The contents of bits 12:63 of srcl are placed into bits 12:63 of VSR[ XT].

The contents of doubleword element 1 of VSR[ XT] are undefined.

\section*{Special Registers Altered:}

None

\section*{Programming Note}

This instruction can be used to produce a single-precision result.

VSR Data Layout for xsiexpdp


VSX Scalar Insert Exponent Quad-Precision X-form
xsiexpqp VRT,VRA,VRB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 63 & VRT & VRA & VRB & & 868 & \(1_{31}\) \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablell
VSR[VRT+32], bit[0] \(\leftarrow\) VSR[VRA+32], bit[0]
VSR[VRT+32], bit[1:15] \&VSR[VRB+32].dword[0], bit[49:63]
VSR[VRT+32], bit[16:127] \& VSR[VRA+32], bit[16:127]
The contents of bit 0 of VSR[VRA+32] are placed into bit 0 of VSR[ VRT +32].

The contents of bit 49:63 of doubleword element 0 of VSR[ VRB +32 ] are placed into bits 1:15 of VSR[ VRT +32].

The contents of bit 16:127 of VSR[VRA+32] are placed into bits 16:127 of VSR[ VRT +32].

Special Registers Altered:
None
\| VSR Data Layout for xsiexpqp
VSR[ VRA+32]


VSR[ VRB+32]
\(\square\)
VSR[VRT+32]
\(\square\)

VSX Scalar Multiply-Add Double-Precision XX3-form
xsmaddadp XT,XA,XB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & 11 & A & 16 & B & 21 & 33 & \(\left\lvert\, \begin{aligned} & \text { Axbx } \\ & 29303031\end{aligned}\right.\) \\
\hline
\end{tabular}
xsmaddmdp XT,XA,XB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & 11 & A & 16 & B & 21 & 41 & \[
\left|\begin{array}{l}
\mid x X B X X T X \\
293031
\end{array}\right|
\] \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline XT & \(\leftarrow \mathrm{TX} \| \mathrm{T}\) \\
\hline XA & \(\leftarrow A X \| A\) \\
\hline ХВ & \(\leftarrow B X|\mid B\) \\
\hline \multicolumn{2}{|l|}{reset_xflags()} \\
\hline src1 & \(\leftarrow \mathrm{VSR}[\mathrm{XA}]\{0: 63\}\) \\
\hline src2 & \(\leftarrow\) "xSmaddadp" ? VSR[XT] \{0:63\} : VSR[XB] \{0:63\} \\
\hline & \(\leftarrow\) "xsmaddadp" ? VSR[XB] \(00: 63\}\) : VSR[XT] \{0:63\} \\
\hline v\{0:inf\} & \(\leftarrow\) MultiplyAddFP(src1, src3,src2) \\
\hline result \(00: 63\}\) & (3) \(\leftarrow\) RoundToDP(RN, v) \\
\hline if(vxsnan_flag) & flag) then SetFX(VXSNAN) \\
\hline if(vximz_flag) & lag) then SetFX(VXIMZ) \\
\hline if(vxisi_flag) & lag) then SetFX(VXISI) \\
\hline if(ox_flag) & then \(\operatorname{SetFX}(0 X)\) \\
\hline if(ux_flag) & then SetFX(UX) \\
\hline if(xx_flag) & then \(\operatorname{SetFX}(x X)\) \\
\hline vex_flag & \(\leftarrow\) VE \& (vxsnan_flag | vximz_flag | vxisi_flag) \\
\hline \multicolumn{2}{|l|}{if( ~vex_flag ) then do} \\
\hline \multicolumn{2}{|l|}{VSR[XT] ヶ result || 0xUUUU_UUUU_UUUU_UUUU} \\
\hline \multicolumn{2}{|l|}{FPRF \(\leftarrow\) ClassDP(result)} \\
\hline \multicolumn{2}{|l|}{FR \(\leftarrow\) inc_flag} \\
\hline \multicolumn{2}{|l|}{FI \(\leftarrow\) xx_flag} \\
\hline \multicolumn{2}{|l|}{end} \\
\hline \multicolumn{2}{|l|}{else do} \\
\hline FR \(\leftarrow\) & \(\leftarrow 0 \mathrm{bo}\) \\
\hline \(\mathrm{FI} \leftarrow\) & \(\leftarrow 0 \mathrm{bO}\) \\
\hline end & \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

For xsmaddadp, do the following.
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
- Let src3 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X B]\).

For xsmaddmdp, do the following.
- Let src2 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X B]\).
- Let src3 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X T]\).
src1 is multiplied \({ }^{[1]}\) by src3, producing a product having unbounded range and precision.

See part 1 of Table 77.
src2 is added \({ }^{[2]}\) to the product, producing a sum having unbounded range and precision.

The sum is normalized \({ }^{[3]}\).
See part 2 of Table 77.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered}
```

VXSNAN VXISI VXIMZ

```

\footnotetext{
1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( \(\mathrm{G}, \mathrm{R}\), and X ) enter into the computation.
3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}

\section*{572}

VSR Data Layout for xsmadd(alm)dp
src1 = VSR[XA]
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
src2 = xsmaddadp ? VSR[XT] : VSR[XB]
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline src3 = xsmaddadp ? VSR[XB] : VSR[XT] \\
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline DP & \multicolumn{2}{c|}{ undefined } \\
\hline 0 & 64
\end{tabular}

\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src3} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow-\) Infinity & \(p \leftarrow-\) Infinity & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& \mathrm{p} \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(\mathrm{p} \leftarrow+\) Zero & \(\mathrm{p} \leftarrow\)-Zero & \(p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{p} \leftarrow+\) Zero & \(\mathrm{p} \leftarrow+\) Zero & \(\mathrm{p} \leftarrow-\) Zero & \(p \leftarrow-\) Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(\mathrm{p} \leftarrow\)-Zero & \(\mathrm{p} \leftarrow\)-Zero & \(\mathrm{p} \leftarrow+\) Zero & \(\mathrm{p} \leftarrow+\) Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src } 3) \\
& v x \text { snan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(\mathrm{p} \leftarrow\)-Zero & \(\mathrm{p} \leftarrow+\) Zero & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(\mathrm{p} \leftarrow \mathrm{SrC1}\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \[
\begin{aligned}
& p \leftarrow \operatorname{src1} \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \mathrm{p} \leftarrow Q(\text { src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q \text { (src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q \text { (src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}

\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src2} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \[
\begin{array}{|l|}
\hline v \leftarrow d Q N a N \\
\text { vxisi_flag } \leftarrow 1
\end{array}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(v \leftarrow A(p, s r c 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow A\left(p, s c^{\prime} 2\right)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{v} \leftarrow \mathrm{SrC2}\) & \(\mathrm{V} \leftarrow\)-Zero & \(v \leftarrow\) Rezd & \(\mathrm{V} \leftarrow \mathrm{src} 2\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{src} 2\) & \[
\begin{array}{|l|}
\hline v \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \(v \leftarrow\) Rezd & \(\mathrm{v} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{src} 2\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(v \leftarrow A(p, s r c 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow A(p, s r c 2)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{dQNaN} \\
& \text { vxisi_flag } \leftarrow 1
\end{aligned}
\] & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \[
\begin{aligned}
& v \leftarrow p \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}

\section*{Explanation:}
\begin{tabular}{|c|c|}
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
\hline src2 & For \(\boldsymbol{x s m a d d a d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT]. For \(\boldsymbol{x s m a d d m d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XB]. \\
\hline src3 & For \(\boldsymbol{x s m a d d a d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XB]. For \(\boldsymbol{x s m a d d m d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT]. \\
\hline dQNaN & Default quiet NaN (0x7FF8_0000_0000_0000). \\
\hline NZF & Nonzero finite number. \\
\hline Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. \\
\hline \(Q(x)\) & Return a QNaN with the payload of x . \\
\hline A(x,y) & Return the normalized sum of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. Note: If \(x=-y, v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline M ( \(\mathrm{x}, \mathrm{y}\) ) & Return the normalized product of floating-point value \(x\) and floating-point value y , having unbounded range and precision. \\
\hline p & The intermediate product having unbounded range and precision. \\
\hline v & The intermediate result having unbounded range and precision. \\
\hline
\end{tabular}

\section*{Table 77.Actions for xsmadd(alm)dp}

\section*{VSX Scalar Multiply-Add Single-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{10}{|l|}{xsmaddasp \(\mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)} \\
\hline \[
06
\] & 6 & T & 11 & A & 16 & B & 21 & 1 &  \\
\hline \multicolumn{2}{|l|}{xsmaddmsp} & & \multicolumn{7}{|l|}{XT, XA, XB} \\
\hline
\end{tabular}

```

reset_xflags()
if "xsmaddasp" then do
src1 \leftarrow VSR[32xAX+A].dword[0]
src2 \leftarrow VSR[32xTX+T].dword[0]
Src3 }\leftarrow\operatorname{VSR[32xBX+B].dword[0]
end
if "xsmaddmsp" then do
src1 \leftarrowVSR[32xAX+A].dword[0]
src2 \leftarrow VSR[32xBX+B].dword[0]
Src3 \leftarrow VSR[32xTX+T].dword[0]
end
v }\leftarrow\mathrm{ MultiplyAddDP(src1,src3,src2)
result \leftarrow RoundTOSP(RN,v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vximz_flag) then SetFX(VXIMZ)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
vex_flag \leftarrow VE \& (vxsnan_flag | vximz_flag | vxisi_flag)
if( ~vex_flag ) then do
VSR[32xTX+T].dword[0] \leftarrow ConvertSPtoSP64(result)
VSR[32xTX+T].dword[1] \leftarrow 0xUUUU_UUUU_UUUU_UUUU
FPRF}\leftarrowClassSP(result
FR }\leftarrow\mathrm{ inc_flag
FI }\leftarrowxx_fla
end
else do
FR}\leftarrow\textrm{ObO
FI}\leftarrow0\textrm{ObO
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

For \(\boldsymbol{x s m a d d a s p}\), do the following.
- Let src1 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X A]\).
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
- Let src3 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X B]\).

For xsmaddmsp, do the following.
- Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
- Let src3 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
src1 is multiplied \({ }^{[1]}\) by src3, producing a product having unbounded range and precision.

See part 1 of Table 78, "Actions for xsmadd(a|m)sp," on page 577.
src2 is added \({ }^{[2]}\) to the product, producing a sum having unbounded range and precision.

The sum is normalized \({ }^{[3]}\).
See part 2 of Table 78, "Actions for xsmadd(a|m)sp," on page 577.

The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\footnotetext{
1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( \(\mathrm{G}, \mathrm{R}\), and X ) enter into the computation.
3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}

\section*{Version 3.0}
```

Special Registers Altered
FPRF FR FI FX OX UX XX
VXSNAN VXISI VXIMZ

```
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{VSR Data Layout for xsmadd(a|m)sp src1 = VSR[XA]} \\
\hline DP & unused \\
\hline \multicolumn{2}{|l|}{src2 = xsmaddasp ? VSR[XT] : VSR[XB]} \\
\hline DP & unused \\
\hline \multicolumn{2}{|l|}{src3 = xsmaddasp ? VSR[XB] : VSR[XT]} \\
\hline DP & unused \\
\hline \multicolumn{2}{|l|}{tgt = VSR[XT]} \\
\hline DP & undefined \\
\hline 0 & \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{\begin{tabular}{l}
Part 1: \\
Multiply
\end{tabular}} & \multicolumn{8}{|c|}{src3} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vxim__flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow-\) Infinity & \(p \leftarrow-\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline -NZF & \(p \leftarrow+\) Infinity & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow+\) Zero & \(p \leftarrow\)-Zero & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow+\) Infinity & \(\mathrm{p} \leftarrow \mathrm{SrC3}\) & \[
\begin{aligned}
& p \leftarrow Q(\operatorname{src} 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline -Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{p} \leftarrow+\) Zero & \(p \leftarrow+\) Zero & \(\mathrm{p} \leftarrow\)-Zero & \(\mathrm{p} \leftarrow\)-Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_lag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \hline p \leftarrow Q(\operatorname{src} 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline - + +Zero & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{p} \leftarrow\)-Zero & \(p \leftarrow\)-Zero & \(\mathrm{p} \leftarrow+\) Zero & \(\mathrm{p} \leftarrow+\) Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_lag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& p \leftarrow Q(\operatorname{src} 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline ¢ +NZF & \(p \leftarrow-\) Infinity & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow\)-Zero & \(p \leftarrow+\) Zero & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& p \leftarrow Q(\operatorname{src} 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline +Infinity & \(p \leftarrow-\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline QNaN & \(\mathrm{p} \leftarrow \mathrm{SrCl}\) & \(\mathrm{p} \leftarrow \mathrm{SrCl}\) & \(p \leftarrow \operatorname{srC1}\) & \(\mathrm{p} \leftarrow \mathrm{SrCl}\) & \(\mathrm{p} \leftarrow \mathrm{SrCl}\) & \(p \leftarrow \operatorname{srCl}\) & \(\mathrm{p} \leftarrow \mathrm{SrCl}\) & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \operatorname{src1} \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline SNaN & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow Q(\operatorname{src} 1) \\
& v x s n a n \_f l a g \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{array}{|l|}
\hline p \leftarrow Q(\text { src1 }) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] & \[
\begin{aligned}
& \hline p \leftarrow Q(\operatorname{src} 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{Q}(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{\begin{tabular}{l}
Part 2: \\
Add
\end{tabular}} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \(V \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow\)-Infinity & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow-\) Infinity & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{dQNaN} \\
\text { vxisi_flag } \leftarrow 1
\end{array}
\] & \(\mathrm{V} \leftarrow \operatorname{src} 2\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] \\
\hline -NZF & \(V \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{A}(\mathrm{p}, \mathrm{src} 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow A(p, s r c 2)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{Src} 2\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline -Zero & \(V \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \(\mathrm{V} \leftarrow-\) Zero & \(\mathrm{v} \leftarrow\) Rezd & \(\mathrm{V} \leftarrow \operatorname{src} 2\) & \(\mathrm{V} \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{Src} 2\) & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline +Zero & \(V \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \(v \leftarrow\) Rezd & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{Src} 2\) & \(\mathrm{V} \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{Src} 2\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline 2.NZF & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{A}(\mathrm{p}, \mathrm{src} 2)\) & \(v \leftarrow p\) & \(V \leftarrow p\) & \(\mathrm{V} \leftarrow \mathrm{A}(\mathrm{p}, \mathrm{src} 2)\) & \(\mathrm{V} \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{Src} 2\) & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow \mathrm{Q} \text { (src2) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline +Infinity & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{dQNaN} \\
& \text { vxisi_flag } \leftarrow 1
\end{aligned}
\] & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(\checkmark \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{Src} 2\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline QNaN \& src1 is a NaN & \(V \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(V \leftarrow p\) & \(V \leftarrow p\) & \(V \leftarrow p\) & \[
\begin{aligned}
& v \leftarrow p \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline QNaN \& src1 not a NaN & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(V \leftarrow p\) & \(\mathrm{V} \leftarrow \operatorname{src} 2\) & \[
\begin{aligned}
& v \leftarrow Q(\operatorname{src} 2) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \multicolumn{9}{|l|}{Explanation:} \\
\hline src1 & \multicolumn{8}{|l|}{The double-precision floating-point value in doubleword element 0 of VSR[XA].} \\
\hline src2 & \multicolumn{8}{|l|}{For \(\boldsymbol{x s m a d d a s p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT]. For \(\boldsymbol{x s m a d d m s p}\), the double-precision floating-point value in doubleword element 0 of VSR[XB].} \\
\hline src3 & \multicolumn{8}{|l|}{For xsmaddasp, the double-precision floating-point value in doubleword element 0 of VSR[XB]. For \(\boldsymbol{x s m a d d m s p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT].} \\
\hline dQNaN & \multicolumn{8}{|l|}{Default quiet NaN (0x7FF8_0000_0000_0000).} \\
\hline NZF & \multicolumn{8}{|l|}{Nonzero finite number.} \\
\hline Rezd & \multicolumn{8}{|l|}{Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands.} \\
\hline \(Q(x)\) & \multicolumn{8}{|l|}{Return a QNaN with the payload of x .} \\
\hline A(x,y) & \multicolumn{8}{|l|}{Return the normalized sum of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. Note: If \(x=-y, v\) is considered to be an exact-zero-difference result (Rezd).} \\
\hline M ( \(\mathrm{x}, \mathrm{y}\) ) & \multicolumn{8}{|l|}{Return the normalized product of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision.} \\
\hline p & \multicolumn{8}{|l|}{The intermediate product having unbounded range and precision.} \\
\hline v & \multicolumn{8}{|l|}{The intermediate result having unbounded range and precision.} \\
\hline
\end{tabular}

Table 78.Actions for \(\operatorname{xsmadd}(a \mid m) s p\)

VSX Scalar Multiply-Add Quad-Precision [using round to Odd] X-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{xsmaddqp xsmaddqpo}} & \multicolumn{3}{|l|}{\multirow[t]{2}{*}{\begin{tabular}{l}
VRT,VRA,VRB \\
VRT,VRA,VRB
\end{tabular}}} & \multicolumn{2}{|r|}{\multirow[t]{2}{*}{\[
\begin{aligned}
& (R O=0) \\
& (R O=1)
\end{aligned}
\]}} \\
\hline & & & & & & \\
\hline \[
063
\] & \[
{ }_{6} \text { VRT }
\] & \({ }_{11}\) VRA & \({ }_{16}\) VRB & 21 & 388 & \begin{tabular}{|r} 
RO \\
31
\end{tabular} \\
\hline \multicolumn{7}{|l|}{if MSR.VSX=0 then VSX_Unavailablell} \\
\hline \multicolumn{7}{|l|}{reset_xflags()} \\
\hline \multicolumn{7}{|l|}{srcl \(\leftarrow\) bfp_CONVERT_FROM_BFP128(VSR[VRA+32])} \\
\hline \multicolumn{7}{|l|}{srct \(\leftarrow\) bfp_CONVERT_FROM_BFP128(VSR[VRT+32])} \\
\hline \multicolumn{7}{|l|}{src3 \(\leftarrow\) bf p_CONVERT_FROM_BPP128(VSR[VRB+32])} \\
\hline \multicolumn{7}{|l|}{\(v \quad \leftarrow\) bfp_MULTIPLY_ADD(srcl , src3, src2)} \\
\hline \multicolumn{7}{|l|}{ind \(\leftarrow\) bfp_ROUND_TO_BFP128(RO, FPSCR, RN, v)} \\
\hline \multicolumn{7}{|l|}{result \(\leftarrow\) bfp_CONVERT_TO_BFP128(rnd)} \\
\hline \multicolumn{7}{|l|}{if(vxsnan_flag) then SetFX(FPSCR. VXSNAN)} \\
\hline \multicolumn{7}{|l|}{if(vximz_flag) then SetFX(FPSCR. VXI MZ)} \\
\hline \multicolumn{7}{|l|}{if(ox_flag) then SetFX(FPSCR. OX)} \\
\hline \multicolumn{7}{|l|}{if(ux_flag) then SetFX(FPSCR. UX)} \\
\hline \multicolumn{7}{|l|}{if(xx_flag) then SetFX(FPSCR, XX)} \\
\hline \multicolumn{7}{|l|}{vx_flag \(\leftarrow\) vxsnan_flag \({ }_{\text {dximz }}\) flag} \\
\hline \multicolumn{7}{|l|}{ex_flag \(\leftarrow\) FPSCR.VE \& Vx_flag} \\
\hline \multicolumn{7}{|l|}{if ex_flag=0 then do} \\
\hline \multicolumn{7}{|l|}{VSR[VRT+32] \(\leftarrow\) result} \\
\hline \multicolumn{7}{|l|}{FPSCR, FPRF \(\leftarrow\) fprf_CLASS_BFP128(result)} \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{FPSCR.FR \(\leftarrow\left(v x_{\text {_ }} f(a g=0)\right.\) \& inc_flag} \\
\hline \multicolumn{7}{|l|}{FPSCR, Fl \(\leftarrow\left(v x_{\sim} f \mid a g=0\right) \& x x_{\sim} f \mid a g\)} \\
\hline
\end{tabular}

Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let srcl be the floating-point value in VSR[VRT+32] represented in quad-precision format.

Let src3 be the floating-point value in VSR[VRB+32] represented in quad-precision format.

If either \(\operatorname{srcl}, \operatorname{src} 2\), or \(\operatorname{srcs}\) is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN is set to 1.

If srcl is an Infinity value and srcs is a Zero value, or if srcl is a Zero value and src3 is an Infinity value, an Invalid Operation exception occurs and VXI MZ is set to 1.

If \(\mathrm{srcl}_{2}\) and the product of srcl and \(\mathrm{srcs}^{2}\) are Infinity values having opposite signs, an Invalid Operation exception occurs and VXISI is set to 1 .

If srCl is a Signalling NaN , the result is the Quiet NaN corresponding to srcl .

Otherwise, if srcl is a Quiet NaN , the result is srcl .
Otherwise, if \(\mathrm{SrCl}_{2}\) is a Signalling NaN , the result is the Quiet NaN corresponding to sic2.

Otherwise, if src 2 is a Quiet NaN , the result is src .
Otherwise, if \(\mathrm{SrCs}^{\mathrm{rc}}\) is a Signalling NaN , the result is the Quiet NaN corresponding to src3.

Otherwise, if \(\operatorname{src} 3\) is a Quiet NaN , the result is src 3 .
Otherwise, if srcl is an Infinity value and srcs is a Zero value, or if srcl is a Zero value and srcs is an Infinity value, the result is the default Quiet \(\mathrm{NaN}^{[1]}\).

Otherwise, if the product of srcl and srcs , and srct are Infinity values having opposite signs, the result is the default Quiet NaN .

Otherwise, do the following.
srcl is multiplied by src3, producing a product having unbounded significand precision and exponent range.

See part 1 of Table 77. "Actions for xsmadd(a|m)dp".
srce is added to the product, producing a sum having unbounded range and precision.

See part 2 of Table 77. "Actions for xsmadd(a|m)dp".

If the intermediate result is Tiny (i.e., the unbiased exponent is less than -16382) and \(U E=0\), the significand is shifted right \(N\) bits, where \(N\) is the difference between - 16382 and the unbiased exponent of the intermediate result. The exponent of the intermediate result is set to the value - 16382.

If \(R O=1\), let the rounding mode be Round to Odd. Otherwise, let the rounding mode be specified by RN. Unless the result is an Infinity or a Zero, the intermediate result is rounded to quad-precision using the specified rounding mode.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into VSR[ VRT +32] in quad-precision format.


FPRF is set to the class and sign of the result. \(F R\) is set to indicate if the rounded result was incremented. Fl is set to indicate the result is inexact.

If a trap-disabled Invalid Operation exception occurs,
FR and FI are set to 0 .
If a trap-enabled Invalid Operation exception occurs,
VSR[VRT+32] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

Special Registers Altered:
FPRF FR FI
FX VXSNAN VXIMZ VXISI OX UX XX




| Table 79.Actions for xsmaddqp[o]

\section*{VSX Scalar Maximum Double-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|}
\hline xsmax & \multicolumn{4}{|c|}{XT, XA, XB} \\
\hline \[
0
\] & \[
{ }_{6} \quad \mathrm{~T}
\] & 11 & 16 & 21 \\
\hline \multicolumn{5}{|l|}{XT \(\quad \leftarrow\) TX \| T} \\
\hline \multicolumn{5}{|l|}{\(X A \quad \leftarrow A X \| A\)} \\
\hline \multicolumn{5}{|l|}{\multirow[t]{2}{*}{\[
\begin{aligned}
& X B \\
& \text { reset_xflags() }
\end{aligned} \leftarrow \mathrm{BX} \| \mathrm{B}
\]}} \\
\hline & & & & \\
\hline \multicolumn{5}{|l|}{src1 \(\quad \leftarrow \operatorname{VSR}[\mathrm{XA}]\{0: 63\}\)} \\
\hline \multicolumn{5}{|l|}{src2 \(\leftarrow \operatorname{VSR}[\mathrm{XB}]\{0: 63\}\)} \\
\hline \multicolumn{5}{|l|}{result \(\{0: 63\} \leftarrow\) MaximumDP( \(\operatorname{src1}\), src2)} \\
\hline \multicolumn{5}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{5}{|l|}{vex_flag \(\leftarrow\) VE \& vxsnan_flag} \\
\hline \multicolumn{5}{|l|}{if( ~vex_flag ) then do} \\
\hline \multicolumn{5}{|r|}{VSR \([\mathrm{XT}] \leftarrow\) result || 0xUUUU_UUUU_UUUU_UUUU} \\
\hline \multicolumn{5}{|l|}{end} \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].

If src1 is greater than src2, src1 is placed into doubleword element 0 of VSR[XT]. Otherwise, src2 is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

The maximum of +0 and -0 is +0 . The maximum of a QNaN and any value is that value. The maximum of any value and an SNaN is that SNaN converted to a QNaN.

FPRF, FR and FI are not modified.
If a trap-enabled invalid operation exception occurs, \(\mathrm{VSR}[\mathrm{XT}]\) is not modified.

See Table 80.

\section*{Special Registers Altered}

FX VXSNAN

\section*{Programming Note}

This instruction can be used to operate on single-precision source operands.

VSR Data Layout for xsmaxdp
src1 \(=\mathrm{VSR}[\mathrm{XA}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline src2 = VSR[XB] & unused \\
\hline DP &
\end{tabular}
tgt \(=\operatorname{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline & \multicolumn{2}{|c|}{ undefined } \\
\hline 0 & 64 \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & T(sicl) & T(sro2) & T(sro2) & T(sro2) & T (scr2) & T (scre) & T(scri) & \(T(Q(s \mathrm{sc} 2))\) fx(VXSNAN) \\
\hline -NZF & T (sicl) & T(M(ssci, ,sc2)) & T (sro2) & T (sro2) & T (scr2) & T (scre2) & T(scri) & \(T\left(Q\left(s r^{2}\right)\right)\) fx(VXSNAN) \\
\hline -Zero & T(srci) & T (stri) & T (scri) & T(str2) & T(str2) & \(T(\) scor \()\) & T (scri) & \(T\left(Q\left(s r^{2}\right)\right)\) fx(VXSNAN) \\
\hline - +Zero & T(srci) & T (stri) & T(scri) & T (scol) & T(str2) & \(T(s \mathrm{c} 2)\) & T (scri) & \(T\left(Q\left(s r^{2}\right)\right)\) tx(VXSNAN) \\
\hline +NZF & T(srci) & T(srci) & T (stri) & T(srci) & T(MMssc1,sc2)) & T(str2) & T(scri) & \(T\left(Q\left(s r_{2}\right)\right)\) tx(VXSNAN) \\
\hline +Infinity & T(srci) & T (stri) & T (stc1) & T (stc1) & T (stc1) & T(stri) & T (scol) & \(T\left(Q\left(s s_{0}\right)\right)\) fx(VXSNAN) \\
\hline QNaN & T(sro2) & T(str2) & T(str2) & T(sroz) & T(str2) & T(str2) & T(ssc1) & T (src1) fx(VXSNAN) \\
\hline SNaN & \(T(Q(\operatorname{sc} 11))\) fxVXSNAN & \(T(Q(s \mathrm{c} 11))\) fx(VXSNAN) & \(T\left(Q\left(s c^{\prime} 1\right)\right)\) txVXSNAN & \(T(Q(s \mathrm{cr} 1))\) tx(XXSNAN & \(T(Q(\) scicl \())\) fx(VXSNAN) & \(T(Q(\) srcil \())\) fx(VXSNAN) & T(Q(scil)) tx(VXSNAN) & \(T(Q(s c 1))\) fx(XXSNAN) \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{Explanation:} \\
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
\hline src2 & The double-precision floating-point value in doubleword element 0 of VSR[XT]. \\
\hline NZF & Nonzero finite number. \\
\hline \(Q(x)\) & Return a QNaN with the payload of x . \\
\hline M ( \(\mathrm{x}, \mathrm{y}\) ) & Return the greater of floating-point value \(x\) and floating-point value y . \\
\hline T(x) & The value x is placed in doubleword element 0 of VSR[ XT ] in double-precision format. \\
\hline & The contents of doubleword element 1 of VSR[XT] are undefined. \\
\hline & FPRF, FR and FI are not modified. \\
\hline \(\mathrm{fx}(\mathrm{x})\) & If x is equal to \(0, F X\) is set to \(1 . \mathrm{x}\) is set to 1 . \\
\hline VXSNAN & Floating-Point Invalid Operation Exception (SNaN) status flag, \(\mathrm{FPSCR}_{\mathrm{VXSNAN}}\). If VE=1, update of VSR[XT] is suppressed. \\
\hline
\end{tabular}

Table 80.Actions for xsmaxdp

VSX Scalar Maximum Type-C Double-Precision XX3-form
xsmaxcdp XT,XA,XB

```

if MSR.VSX=O then VSX_Unavailablel)
srcl \leftarrow bfp_CONVERT_FROM_BFP64(VSR[ 32xAX+A].dword[0])
sfC2 \leftarrow bfp_CONVERT_fROM_ BFP64(VSR[32xBX+B],dword[0])
vxsnan_flag \leftarrow(srcl.class="SNaN") | (src2.class="SNaN")
if (src1.type="SNaN") | (src1.type="ONaN")
(scc2.type="SNaN") | (scc2.type="QNaN") then
result }\leftarrowVSR[32\timesBX+B].dword[0
else if bfp_COMPare_gT(srcl,src2) then
result }\leftarrow\mathrm{ VSR[32xAX+A].dword[0]
else
result \leftarrowVSR[32xBX+B].dword[0]
vex_flag \leftarrowFPSCR.VE\& \& venan_flag
if (vxsnan_flag=1) then SetFX(vXSNaN)
if (vex_flag=0) then do
VSR[32xTX+T].dword[0]}\leftarrow\mathrm{ result
VSR[32xTX+T].dword[1] \leftarrow OxUUUO_UUUU_UUUU_UUUU
end

```

I
_ Programming Note
xsmaxcdp can be used to implement the C/C++/Java conditional operation ( \(x>y\) ) ? \(x\) : y for single-precision and double-precision arguments.
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & T(sic2) & T(sic2) & T(sic2) & T(sic2) & T(sic2) & T(sic2) & T(sic2) & \[
\begin{gathered}
\hline \text { T(sic2) } \\
\text { fx(UXSNAN) }
\end{gathered}
\] \\
\hline -NZF & T(sicl) & T(M) siclesic2) & T(sic2) & T(sic2) & T(sic2) & T(sic2) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic}(2) \\
\mathrm{fx}(\mathrm{VXSNAN}) \\
\hline
\end{gathered}
\] \\
\hline -Zero & T(sicl) & T(scri) & T(scr 2) & T(scr 2) & T(sic2) & T(sic2) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s}(\mathrm{C} 2) \\
\mathrm{fx}(\mathrm{VXSNAN}) \\
\hline
\end{gathered}
\] \\
\hline - +Zero & T/sic1) & T/sicl) & T(sic2) & T(sic2) & T(sic2) & T(sic2) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s}(2) 2) \\
\mathrm{fx}(\mathrm{VXSNAN})
\end{gathered}
\] \\
\hline +NZF & T(sicl) & T/sicl) & T/sicl) & T(sicl) & T/MSsic1, sic 2 ) & T(sic2) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s}(\mathrm{C} 2) \\
\mathrm{fx}(\mathrm{XXSNAN})
\end{gathered}
\] \\
\hline +Infinity & T(sicl) & T(sicl) & T(sicl) & T(sicl) & T(scri) & T(sic2) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s}(\mathrm{C} 2) \\
\mathrm{fx}(\mathrm{UXSNAN}) \\
\hline
\end{gathered}
\] \\
\hline QNaN & T(sic2) & T(ssc2) & T(sic2) & T(sic2) & T(scr 2 ) & T(sic2) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\mathrm{src} 2) \\
\mathrm{fx}(\mathrm{VXSNAN})
\end{gathered}
\] \\
\hline SNaN & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s}(\mathrm{C} 2) \\
\mathrm{fx}(\mathrm{VXSNAN})
\end{gathered}
\] & \[
\begin{gathered}
\hline \text { I(sic2) } \\
\mathrm{fx}(\mathrm{VXSNAN}) \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s}(\mathrm{C} 2) \\
\mathrm{fx}(\mathrm{UXSNAN})
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic}(2) \\
\mathrm{fx}(\mathrm{VXSNAN})
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s} \subset(2) \\
\mathrm{fx}(\mathrm{XXSNAN}) \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s} \subset(2) \\
\mathrm{fx}(\mathrm{XXSNAN})
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s}(\mathrm{C} 2) \\
\mathrm{fx}(\mathrm{VXSNAN})
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\operatorname{ssc}(2) \\
\mathrm{fx}(\mathrm{VXSNAN})
\end{gathered}
\] \\
\hline
\end{tabular}

\section*{Explanation:}
\begin{tabular}{|c|c|}
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[ XA\(]\). \\
\hline src2 & The double-precision floating-point value in doubleword element 0 of VSR[ XT ] . \\
\hline NZF & Nonzero finite number. \\
\hline M ( \(x, y\) ) & Return the greater of floating-point value x and floating-point value y . \\
\hline \(T(x)\) & The value \(x\) is placed in doubleword element 0 of VSR[ PT] \(^{\text {a }}\) in double-precision format. \\
\hline & The contents of doubleword element 1 of VSR[ \(X T]\) are undefined. \\
\hline & FPRF, FR and FI are not modified. \\
\hline \(\mathrm{fx}(\mathrm{x})\) & If X is equal to \(0, F \mathrm{X}\) is set to \(1 . \mathrm{x}\) is set to 1 . \\
\hline VXSNAN & Floating-Point Invalid Operation Exception (SNaN) status flag, VXS NAN. If VE=1, update of VSR[ XT ] is suppressed. \\
\hline
\end{tabular}

\section*{| Table 81.Actions for xsmaxcdp}

VSX Scalar Maximum Type-J Double-Precision XX3-form
xsmaxjdp XT,XA,XB

```

if MSR.VSX=O then VSX_Unavailablel)
srcl \leftarrow bf p_CONVERT_FROM_BFP64(VSR[32xAX+A],dword[0])
srC2 \leftarrow bfp_CONVERT_FROM_BFP64(VSR[32xBX+B].dword[0])
vxsnan_flag \leftarrow(srcl.class="SNaN")|(src2.class="SNaN")
if (srcl.type="SNaN") | (srcl,type="ONaN") then
result }\leftarrow\mathrm{ VSR[32xAX+A].dword[0]
else if (src2.type="SNaN")|(sic2.type="QNaN") then
result }\leftarrowV\mp@subsup{V}{R}{[}[32\timesBX+B].dword[0
else if (src1.type="Zero") \& (src2.type="Zero") then
if (srcl.sign=0) | (src2.sign=0) then
result \leftarrow Ox0000_0000_0000_0000 11 +Zero
else
result \leftarrow0x8000_0000_0000_0000 1/ -Zero
else if bfp_COMPARE_GT(srcl,sic2) then
result \leftarrow VSR[32xAX+A].dword[0]
else
result }\leftarrow\mathrm{ VSR[32xBX+B].dword[0]
vex_flag \leftarrowFPSCR.VE\& \&xsnan_flag
if (vxsnan_flag=1) then SetFX(FPSCR. UXSNAN)
if(vex_flag=0) then do
VSR[ [32xTX+T].dword[0] \leftarrowbfp64_CONVERT_from_bfp(result)
VSR[32xTX+T],dword[1] \leftarrow OxUUUU_UUUU_UUUU_UUUU
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let srcl be the double-precision floating-point value in doubleword 0 of VSR[ XA].

Let src2 be the double-precision floating-point value in doubleword 0 of VSR[ XB].

If srcl or srcl is a SNaN, an Invalid Operation exception occurs.

If srcl is a NaN, result is srcl.
Otherwise, if \(\mathrm{src}_{2}\) is a NaN , result is \(\mathrm{src2}\).
Otherwise, if \(\operatorname{srcl}\) is a Zero and \(\mathrm{srcl}^{2}\) is a Zero and either srcl or srcl is a +Zero, the result is +Zero.

Otherwise, if srcl is a -Zero and srcl is a -Zero, the result is -Zero.

Otherwise, if \(\mathrm{srcl}_{1}\) is greater than \(\mathrm{src2}\), result is \(\operatorname{srcl}\).
Otherwise, result is src2.
The contents of doubleword 0 of VSR[ XT] are set to the value result.

The contents of doubleword 1 of VSR[XT] are undefined.

If a trap-enabled Invalid Operation occurs, VSR[XT] is not modified.

\section*{Special Registers Altered:}

FX VXSNAN

Version 3.0

I
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & T( \(\cdot 1 \mathrm{NF}\) ) & T(src2) & T(sic2) & T(src2) & T(src2) & T(src2) & T(src2) & \[
\begin{gathered}
\text { T(sic2) } \\
\text { fx(VXSNAN) } \\
\hline
\end{gathered}
\] \\
\hline -NZF & T(srci) & T(M(srcl, src2) & T(sic2) & T(src2) & T(src2) & T(src2) & T(src2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\mathrm{fx}(\mathrm{VXSNAN})
\end{gathered}
\] \\
\hline -Zero & T(srcl) & T(srch) & T(-Zerol & T + +2erol & T(src2) & T(src2) & T(src2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\text { fx(VXSNAN) }
\end{gathered}
\] \\
\hline \(\overline{\text { ¢ }}\) +Zero & T(srci) & T(sicl) & T( + Zerol & T( +Zerol & T(src2) & T(sic2) & T(src2) & \[
\begin{gathered}
\text { T(sic2) } \\
\text { fx(VXSNAN) } \\
\hline
\end{gathered}
\] \\
\hline ¢ +NZF & T(srci) & T/srci) & T/sricl) & T(srci) & T/M(srcl, sicl \({ }^{\text {a }}\) & T(src2) & T(src2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\mathrm{fx}(\text { VXSNAN })
\end{gathered}
\] \\
\hline +Infinity & T(srcl) & T(srch) & T(sricl) & T(srcl) & T(srcl) & T ( + WF) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\text { fx(VXSNAN) }
\end{gathered}
\] \\
\hline QNaN & T(srci) & T(srci) & T(srcl) & T(srci) & T(srcl) & T(srcl) & T(srch) & \[
\begin{gathered}
\hline \text { T(sicl) } \\
\text { fx(VXSNAN) } \\
\hline
\end{gathered}
\] \\
\hline SNaN & \[
\begin{gathered}
\mathrm{T}(\text { sicl) } \\
\mathrm{fx}(\text { VXSNAN })
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\mathrm{sic} 1) \\
\mathrm{fx}(\text { VXSNAN })
\end{gathered}
\] & \[
\begin{gathered}
\text { T(sic1) } \\
\text { fX(VXSNAN) }
\end{gathered}
\] & \[
\begin{gathered}
\text { T(sici) } \\
\text { fx(VXSNAN) }
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\text { srci) } \\
\mathrm{fx}(\text { VXSNAN })
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\text { srci) } \\
\mathrm{fx}(\text { UXSNAN) }
\end{gathered}
\] & \[
\begin{gathered}
\text { T(sicl) } \\
\text { fx(VXSNAN) }
\end{gathered}
\] & \[
\begin{gathered}
\text { T(sicl) } \\
\text { fX(VXSNAN) }
\end{gathered}
\] \\
\hline
\end{tabular}

\section*{Explanation:}
\begin{tabular}{|c|c|}
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[ XA\(]\). \\
\hline src2 & The double-precision floating-point value in doubleword element 0 of VSR[ XT ] . \\
\hline NZF & Nonzero finite number. \\
\hline M ( \(x, y\) ) & Return the greater of floating-point value \(x\) and floating-point value y . \\
\hline \(T(x)\) & The value \(x\) is placed in doubleword element 0 of VSR[ \(X T]\) in double-precision format. \\
\hline & The contents of doubleword element 1 of VSR[ XT\(]\) are undefined. \\
\hline & FPRF,FR and FI are not modified. \\
\hline \(\mathrm{fx}(\mathrm{x})\) & If X is equal to \(0, F \mathrm{X}\) is set to \(1 . \mathrm{X}\) is set to 1 . \\
\hline VXSNAN & Floating-Point Invalid Operation Exception (SNaN) status flag, VXSNAN. If VE=1, update of VSR[ XT\(]\) is suppressed. \\
\hline
\end{tabular}

\section*{I Table 82.Actions for xsmaxjdp}

\section*{VSX Scalar Minimum Double-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|l|}{xsmindp \(\quad\) XT, XA, XB} \\
\hline \[
\begin{array}{ll} 
& 60 \\
0
\end{array}
\] & \[
{ }_{6} \quad \mathrm{~T}
\] & \[
{ }_{11} A
\] & 116 & 21 & 168 & \(|\)\begin{tabular}{c|c|c|c|} 
AX \\
29 & TX \\
29 & 30 & 31
\end{tabular} \\
\hline \multicolumn{7}{|l|}{\(X T \quad \leftarrow T X \| T\)} \\
\hline \multicolumn{7}{|l|}{\(X A \quad \leftarrow A X \| A\)} \\
\hline \multicolumn{7}{|l|}{} \\
\hline \multicolumn{7}{|l|}{reset_xflags()} \\
\hline \multicolumn{7}{|l|}{src1 \(\quad \leftarrow \operatorname{VSR}[\mathrm{XA}]\{0: 63\}\)} \\
\hline \multicolumn{7}{|l|}{src2 \(\leftarrow\) VSR[XB]\{0:63\}} \\
\hline \multicolumn{7}{|l|}{result \(\{0: 63\} \leftarrow\) MinimumDP( \(\operatorname{src1}\), src2)} \\
\hline \multicolumn{7}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{7}{|l|}{vex_flag \(\leftarrow\) VE \& vxsnan_flag} \\
\hline \multicolumn{7}{|l|}{if( ~vex_flag ) then do} \\
\hline \multicolumn{7}{|l|}{\multirow[t]{2}{*}{VSR \([X T] \leftarrow\) result \(\|\) 0xUUUU_UUUU_UUUU_UUUU
end}} \\
\hline & & & & & & \\
\hline
\end{tabular}

VSR Data Layout for xsmindp
srcl = VSR[XA]
\begin{tabular}{|c|c|}
\hline \multicolumn{1}{|c|}{ DP } & unused \\
\hline srC2 \(=\) VSR[ XB] & unused \\
\hline \multicolumn{2}{|c|}{ DP } \\
\hline t gt \(=\) VSR[ XT] & \\
\hline \multicolumn{2}{|c|}{ DP } \\
\hline 0 & undefined \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let srcl be the double-precision floating-point value in doubleword element 0 of VSR[ XA] .

Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

If srcl is less than srcl , srcl is placed into doubleword element 0 of VSR[XT] in double-precision format. Otherwise, src2 is placed into doubleword element 0 of VSR[ XT] in double-precision format.

The contents of doubleword element 1 of VSR[ XT] are undefined.

The minimum of +0 and -0 is -0 . The minimum of a QNaN and any value is that value. The minimum of any value and an SNaN is that SNaN converted to a QNaN.

FPRF, FR and FI are not modified.
If a trap-enabled invalid operation exception occurs, VSR[ XT] is not modified.

See Table 83.

\section*{Special Registers Altered}

FX VXSNAN

\section*{Programming Note}

This instruction can be used to operate on single-precision source operands.
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & T(sicli) & T(srci) & T(scri) & T(scri) & T(stri) & T(scri) & T(scri) & \(\mathrm{T}(Q \mathrm{Qsc} 2))\) fx(VXSNAN) \\
\hline -NZF & T(sro2) & T(M(ssci, ssc2)) & T (scri) & T (stc1) & T (scri) & T(scri) & T (scri) & \(T(Q(\) scr2)) fx(VXSNAN) \\
\hline -Zero & T(sro2) & T(str2) & T(srci) & T (scri) & T (scol) & \(T\) (scri) & \(T\) (scri) & \(T(Q(\) scr2 ) \()\) tx(VXSNAN) \\
\hline - +Zero & T(sro2) & T(str2) & T (str2) & T (sc1) & T (scri) & \(T\) (scri) & T (scri) & \(T(Q(s c 2))\) fx(VXSNAN) \\
\hline +NZF & T(sroz) & T(sroz) & T(str2) & T(str2) & T(MMssc1,sc2)) & T(scri) & \(T\) (scri) & \(T(Q(s c 2))\) fx(VXSNAN) \\
\hline +Infinity & T(srce) & T(sro2) & T(str2) & \(T\) (str2) & T(stre) & T(scri) & \(T\) (scri) & \(T\left(Q\left(s r_{2}\right)\right)\) fx(VXSNAN) \\
\hline QNaN & T(stre) & T(str2) & T(str2) & T(str2) & T(sroz) & T(str2) & T (scol) & \[
\begin{gathered}
\mathrm{T}(\mathrm{src} 1) \\
\text { fx(VXSNAN) }
\end{gathered}
\] \\
\hline SNaN & \(T(Q(\operatorname{sc} 11))\) tx(XXSNAN & \(T(Q(\) scci 1\())\) fx(VXSNAN) & \(T(Q(\) srci 1\())\) fx(VXSNAN & \(T(Q(\operatorname{src} 1))\) fx(VXSNAN & \(T(Q(\) scri) ) fx(VXSNAN) & \[
T(Q(\operatorname{src} 1))
\]
fx(VXSNAN) & \(T(Q(\) srci 1\())\) fx(VXSNAN & \(T(Q(\) srci1 \())\) fx(VXSNAN) \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{Explanation:} \\
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
\hline src2 & The double-precision floating-point value in doubleword element 0 of VSR[XT]. \\
\hline NZF & Nonzero finite number. \\
\hline \(Q(x)\) & Return a QNaN with the payload of x . \\
\hline M ( \(\mathrm{x}, \mathrm{y}\) ) & Return the lesser of floating-point value \(x\) and floating-point value y . \\
\hline T(x) & The value \(x\) is placed in doubleword element \(i(i \in\{0,1\})\) of VSR[XT] in double-precision format. \\
\hline & The contents of doubleword element 1 of VSR[XT] are undefined. \\
\hline & FPRF, FR and FI are not modified. \\
\hline \(\mathrm{fx}(\mathrm{x})\) & If \(x\) is equal to \(0, F X\) is set to \(1 . x\) is set to 1 . \\
\hline VXSNAN & Floating-Point Invalid Operation Exception (SNaN) status flag, FPSCR \({ }_{\text {VxSNAN }}\). If VE=1, update of VSR[XT] is suppressed. \\
\hline
\end{tabular}

Table 83.Actions for xvmindp

\section*{VSX Scalar Minimum Type-C Double-Precision XX3-form}
Xsmincdp \(\mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)
\begin{tabular}{|l|l|l|l|l|l|l|l|l|}
\hline 60 & T \\
\hline 0 & & 6 & & 11 & A & & B & 16 \\
\hline
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable()
srcl \leftarrow bfp_CONVERT_FROM_BFP64(VSR[ 32xAX+A].dword[0])
sfC2 \leftarrow bfp_CONVERT_fROM_ BFP64(VSR[32xBX+B],dword[0])
vxsnan_flag \leftarrow(srcl.class="SNaN") | (src2.class="SNaN")
if (src1.type="SNaN") | (src1.type="ONaN")
(scc2.type="SNaN") | (scc2.type="QNaN") then
result }\leftarrow\mathrm{ VSR[ [2 ABX +B].dword[0]
else if bfp_ COMPare_Lt(sicl,src2) then
result }\leftarrow\mathrm{ VSR[32xAX+A].dword[0]
else
result }\leftarrow\mathrm{ VSR[ [2xBX+B].dword[0]
vex_flag \leftarrowFPSCR.VE \& vxSnan_flag
if (vxsnan_flag=1) then SetFX(VXSNAN)
if (vex_flag=0) then do
VSR[32xTX+T].dword[0]}\leftarrow\mathrm{ result
VSR[32xTX+T].dword[1] \leftarrow OxUUUZ_UUUU_UUUU_UUUU
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let srcl be the double-precision floating-point value in doubleword 0 of VSR[ XA].

Let src2 be the double-precision floating-point value in doubleword 0 of VSR[ XB].

If srcl or srcl is a SNaN, an Invalid Operation exception occurs.

If eithersrcl orsrcl is a NaN, result is srcl.
Otherwise, if srcl is less than \(\mathrm{srch}^{2}\), result is srcl .
Otherwise, result is src2.
The contents of doubleword 0 of VSR[ XT] are set to the value result.

The contents of doubleword 1 of VSR[XT] are undefined.

If a trap-enabled Invalid Operation occurs, VSR[ XT] is not modified.

\section*{Special Registers Altered:}

FX VXSNAN

\section*{Programming Note}
xsmincdp can be used to implement the C/C++/Java conditional operator \((x<y) ? x: y\) for single-precision and double-precision arguments.

Version 3.0

I
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & T(sic2) & T(sic1) & T(sicl) & T(sicl) & T(sic1) & T(sicl) & T(sic2) & \[
\begin{gathered}
\hline \text { I(sic2) } \\
f \times(\text { (XXSNAN) }
\end{gathered}
\] \\
\hline -NZF & T(sic2) & T(M) sicl, sic2) & T(sic1) & T(sicl) & T(sic1) & T(sicl) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic}(2) \\
\mathrm{fx}(\mathrm{VXSNAN}) \\
\hline
\end{gathered}
\] \\
\hline -Zero & T(sic2) & T(ssc2) & T(ssc2) & T(ssc2) & T(sicl) & T(sic1) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s}(\mathrm{c} 2) \\
\mathrm{fx}(\mathrm{VXSNAN}) \\
\hline
\end{gathered}
\] \\
\hline - +Zero & T(sic2) & T/sic2) & T(sic2) & T(scr 2 ) & T(sicl) & T(sic1) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s}(\mathrm{C} 2) \\
\mathrm{fx}(\mathrm{VXSNAN}) \\
\hline
\end{gathered}
\] \\
\hline - & T(sic2) & T/sic2) & T(sic2) & T(sic2) & T(M|sscl, sic 2 ) & T(sic1) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic}(2) \\
\mathrm{fx}(\mathrm{VXSNAN}) \\
\hline
\end{gathered}
\] \\
\hline +Infinity & T(sic2) & T/sic2) & T/sic2) & T(sic2) & T/sic2) & T(ssc2) & T(sic2) & \[
\begin{gathered}
\hline \mathrm{T}(\mathrm{~s}(\mathrm{C} 2) \\
\mathrm{fx}(\mathrm{XXSNAN}) \\
\hline
\end{gathered}
\] \\
\hline QNaN & T(sic2) & T/sic2) & T/sic2) & T/sic2) & T/sic2) & T(ssc2) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{ssc}(2) \\
\mathrm{fx}(\mathrm{VXSNAN}) \\
\hline
\end{gathered}
\] \\
\hline SNaN & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s} \text { (C2) } \\
\mathrm{fx}(\mathrm{XXSNAN}) \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{I}(\mathrm{~s} \subset \mathrm{C}) \\
\mathrm{fx}(\mathrm{yXsNAN}) \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\mathrm{fx}(\mathrm{VXSNAN}) \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s},(2) \\
\mathrm{fx}(\mathrm{XXSNAN}) \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{I}(\mathrm{~s} \subset \mathrm{C}) \\
\mathrm{fx}(\mathrm{XXSNAN}) \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s},(2) \\
\mathrm{fx}(\mathrm{XXSNAN}) \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\mathrm{fx}(\mathrm{VXSNAN}) \\
\hline
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\mathrm{~s}(\mathrm{C} 2) \\
\mathrm{fx}(\mathrm{VXSNAN})
\end{gathered}
\] \\
\hline
\end{tabular}

\section*{Explanation:}
\begin{tabular}{|c|c|}
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[ XA ] . \\
\hline srct & The double-precision floating-point value in doubleword element 0 of VSR[ XT ] . \\
\hline NZF & Nonzero finite number. \\
\hline M ( \(x, y\) ) & Return the lesser of floating-point value \(x\) and floating-point value \(y\). \\
\hline T( x ) & The value \(x\) is placed in doubleword element 0 of VSR[ \(X T]\) in double-precision format. The contents of doubleword element 1 of VSR[ \(X T\) ] are undefined. FPRF, FR and FI are not modified. \\
\hline \(\mathrm{fx}(\mathrm{x})\) & If X is equal to \(0, F \mathrm{~F}\) is set to \(1 . \mathrm{x}\) is set to 1 . \\
\hline VXSNAN & Floating-Point Invalid Operation Exception (SNaN) status flag, VXSNAN. If \(\mathrm{VE}=1\), update of VSR[ XT\(]\) is suppressed. \\
\hline
\end{tabular}

\section*{| Table 84.Actions for xsmincdp}

\section*{VSX Scalar Minimum Type-J Double-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{xsminjdp \(\quad X T, X A, X B\)} \\
\hline \[
06
\] & \[
6 \quad \text { T }
\] & \[
{ }_{11} A
\] & \[
{ }_{16} \text { B }
\] & \[
2152
\] & \[
\begin{array}{|l|l|l|l|}
\hline A x|B x| T x \\
29 & 30 & 31 \\
\hline
\end{array}
\] \\
\hline
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailablel)
srcl \leftarrow bf p_CONVERT_FROM_BFP64(VSR[32xAX+A],dword[0])
src2 \leftarrow bfp_CONVERT_FROM_BFP64(VSR[32xBX+B].dword[0])
vxsnan_flag \leftarrow(srcl.type="SNaN")|(src2.type="SNaN")
if (srcl,type="SNaN")|(srcl,type="ONaN") then
result }\leftarrow\mathrm{ VSR[32xAX+A],dword[0]
else if (src2.type="SNaN")|(src2.type="QNaN") then
result }\leftarrowV\R[32\timesBX+B].dword[0
else if (src1.type="Zero") \& (src2.type="Zero") then
if (src1.sign=1) | (sic2.sign=1) then
result \leftarrow0x8000_0000_0000_0000 || -Zero
else
result }\leftarrow0x0000_0000_0000_0000 || +Zero
else if bf _COMPARE_LT(srcl,src2) then? srcl: src2
result \leftarrow VSR[32xAX+A].dword[0]
else
result }\leftarrow\mathrm{ VSR[32xBX+B].dword[0]
if (vxsnan_flag=1) then SetFX(FPSCR.VXSNAN)
vex_flag \leftarrow FPSCR.VE \& vxsnan_flag
if(vex_flag=0) then do
VSR[32xTX+T].dword[0] \leftarrowresult
VSR[32xTX+T],dword[1] \leftarrow OxUUUU_UUUU_UUUU_UUUU
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(\mathrm{srcl}^{\mathrm{cl}}\) be the double-precision floating-point value in doubleword 0 of VSR[ XA].

Let src2 be the double-precision floating-point value in doubleword 0 of VSR[ XB].

If srcl or srcl is a SNaN, an Invalid Operation exception occurs.

If srcl is a NaN, result is srcl.
Otherwise, if \(\mathrm{srCl}_{2}\) is a NaN, result is \(\mathrm{src2}\).
Otherwise, if \(\operatorname{srcl}\) is a Zero and \(\mathrm{srcl}^{2}\) is a Zero and eithersrcl orsrc2 is a-Zero, the result is -Zero.

Otherwise, if \(\mathrm{srcl}^{\text {ch }}\) is +Zero and \(\mathrm{srcl}^{2}\) is a +Zero, the result is +Zero.

Otherwise, if srcl is less than \(\mathrm{src2}\), result is \(s \mathrm{cc}\).
Otherwise, result is \(\operatorname{src}\).
The contents of doubleword 0 of VSR[ XT] are set to the value result.

The contents of doubleword 1 of VSR[XT] are undefined.

If a trap-enabled Invalid Operation occurs, VSR[XT] is not modified.

\section*{Special Registers Altered:}

FX VXSNAN

Version 3.0

I
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & T( \(\cdot 1 \mathrm{NF}\) ) & T(srcl) & T/sricl) & T(srci) & T(srci) & T(srci) & T(src2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\mathrm{fx}(\text { VXSNAN })
\end{gathered}
\] \\
\hline -NZF & T(src2) & T(M(srcl, src2) & T(srch) & T(srcl) & T(srcl) & T(srci) & T(src2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\mathrm{fx}(\mathrm{VXSNAN})
\end{gathered}
\] \\
\hline -Zero & T(src2) & T(sic2) & T(-Zero) & T(-Zero) & T(srcl) & T(srcl) & T(src2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\text { fx(VXSNAN) }
\end{gathered}
\] \\
\hline \(\overline{\text { ¢ }}\) +Zero & T(src2) & T(sic2) & T(-Zerol & T( +Zerol & T(srci) & T(srci) & T(src2) & \[
\begin{gathered}
\text { T(sic2) } \\
\text { fx(VXSNAN) } \\
\hline
\end{gathered}
\] \\
\hline ¢ +NZF & T(src2) & T(src2) & T(sic2) & T(src2) & T/M(srcl, sicl \({ }^{\text {a }}\) & T(srch) & T(src2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\mathrm{fx}(\text { VXSNAN })
\end{gathered}
\] \\
\hline +Infinity & T(src2) & T(src2) & T(src2) & T(src2) & T(src2) & T ( + WF) & T(sic2) & \[
\begin{gathered}
\mathrm{T}(\operatorname{sic} 2) \\
\text { fx(VXSNAN) }
\end{gathered}
\] \\
\hline QNaN & T(srci) & T(srci) & T(srcl) & T(srci) & T(srcl) & T(srcl) & T(srch) & \[
\begin{gathered}
\hline \text { T(sicl) } \\
\text { fx(VXSNAN) } \\
\hline
\end{gathered}
\] \\
\hline SNaN & \[
\begin{gathered}
\mathrm{T}(\text { sicl) } \\
\mathrm{fx}(\text { VXSNAN })
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\mathrm{sic} 1) \\
\mathrm{fx}(\text { VXSNAN })
\end{gathered}
\] & \[
\begin{gathered}
\text { T(sic1) } \\
\text { fX(VXSNAN) }
\end{gathered}
\] & \[
\begin{gathered}
\text { T(sici) } \\
\text { fx(VXSNAN) }
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\text { srci) } \\
\mathrm{fx}(\text { VXSNAN })
\end{gathered}
\] & \[
\begin{gathered}
\mathrm{T}(\text { srci) } \\
\mathrm{fx}(\text { UXSNAN) }
\end{gathered}
\] & \[
\begin{gathered}
\text { T(sicl) } \\
\text { fx(VXSNAN) }
\end{gathered}
\] & \[
\begin{gathered}
\text { T(sicl) } \\
\text { fX(VXSNAN) }
\end{gathered}
\] \\
\hline
\end{tabular}

\section*{Explanation:}
\begin{tabular}{|c|c|}
\hline srcl & The double-precision floating-point value in doubleword element 0 of VSR[ XA ] . \\
\hline src2 & The double-precision floating-point value in doubleword element 0 of VSR[ XT ]. \\
\hline NZF & Nonzero finite number. \\
\hline \(M(x, y)\) & Return the greater of floating-point value \(x\) and floating-point value y . \\
\hline T(x) & The value \(x\) is placed in doubleword element 0 of VSR[ PT] \(^{\text {a }}\) in double-precision format. \\
\hline & The contents of doubleword element 1 of VSR[ \(X T]\) are undefined. \\
\hline & FPRF,FR and FI are not modified. \\
\hline \(\mathrm{fx}(\mathrm{x})\) & If X is equal to \(0, F \mathrm{X}\) is set to \(1 . \mathrm{X}\) is set to 1 . \\
\hline VXSNAN & Floating-Point Invalid Operation Exception (SNaN) status flag, VXSNAN. If VE=1, update of VSR[ XT ] is suppressed. \\
\hline
\end{tabular}

I Table 85.Actions for xsminjdp

\section*{VSX Scalar Multiply-Subtract Double-Precision XX3-form}
xsmsubadp \(\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline 06 & 6 & 11 & A & & B & 21 & 49 & \(\left\lvert\, \begin{aligned} & \text { axbxx } \\ & 29303031\end{aligned}\right.\) \\
\hline \multicolumn{2}{|l|}{xsmsubmdp} & \multicolumn{7}{|l|}{XT, XA, XB} \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \({ }^{6} 6\) & 6 & T & 11 & A & 16 & B & 21 & 57 & |AxBx|TX \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline XT & \(\leftarrow T X|\mid T\) \\
\hline XA & \(\leftarrow A X \| A\) \\
\hline ХВ & \(\leftarrow B X|\mid B\) \\
\hline \multicolumn{2}{|l|}{reset_xflags()} \\
\hline src1 & \(\leftarrow \operatorname{VSR}[\mathrm{XA}]\{0: 63\}\) \\
\hline src2 & \(\leftarrow \operatorname{VSR}[\mathrm{XT}]\{0: 63\}\) \\
\hline src3 & \(\leftarrow \mathrm{VSR}[\mathrm{XB}]\{0: 63\}\) \\
\hline src2 & \(\leftarrow\) "xsmsubadp" ? VSR[XT]\{0:63\} : VSR[XB]\{0:63\} \\
\hline & \(\leftarrow\) "xsmsubadp" ? VSR[XB]\{0:63\} : VSR[XT] \{0:63\} \\
\hline v\{0:inf\} & \(\leftarrow\) MultiplyAddDP(src1, src3, NegateDP(src2)) \\
\hline result \(\{0: 63\}\) & 3 \(\leqslant\) RoundToDP( \(\mathrm{RN}, \mathrm{v}\) ) \\
\hline if(vxsnan_fl & flag) then SetFX(VXSNAN) \\
\hline if(vximz_flay) & lag) then SetFX(VXIMZ) \\
\hline if(vxisi_flay & lag) then SetEX(VXISI) \\
\hline if(ox_flag) & then \(\operatorname{SetFX}(0 X)\) \\
\hline if(ux_flag) & then SetFX(UX) \\
\hline if(xx_flag) & then \(\operatorname{SetFX}(X X)\) \\
\hline vex_flag & \(\leftarrow\) VE \& (vxsnan_flag | vximz_flag | vxisi_flag) \\
\hline \multicolumn{2}{|l|}{if( ~vex_flag ) then do} \\
\hline \multicolumn{2}{|l|}{VSR[XT] ヶ result || 0xUUUU_UUUU_UUUU_UUUU} \\
\hline \multicolumn{2}{|l|}{FPRF \(\leftarrow\) ClassDP(result)} \\
\hline \multicolumn{2}{|l|}{FR \(\leftarrow\) inc_flag} \\
\hline \multicolumn{2}{|l|}{FI \(\leftarrow\) xx_flag} \\
\hline \multicolumn{2}{|l|}{end} \\
\hline \multicolumn{2}{|l|}{else do} \\
\hline FR \(\leftarrow\) & \(\leftarrow\) 0b0 \\
\hline \(\mathrm{FI} \leftarrow\) & \(\leftarrow 0 \mathrm{~b} 0\) \\
\hline end & \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
For xsmsubadp, do the following.
- Let src1 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X A]\).
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
- Let src3 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X B]\).

For xsmsubmdp, do the following.
- Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
- Let src3 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
src1 is multiplied \({ }^{[1]}\) by src3, producing a product having unbounded range and precision.

See part 1 of Table 86.
src2 is negated and added \({ }^{[2]}\) to the product, producing a sum having unbounded range and precision.

The result, having unbounded range and precision, is normalized \({ }^{[3]}\).

See part 2 of Table 86.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. FI is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, \(\mathrm{VSR}[\mathrm{XT}]\) and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered}
\[
\begin{aligned}
& \text { FPRF FR FI FX OX UX XX } \\
& \text { VXSNAN VXISI VXIMZ }
\end{aligned}
\]

\footnotetext{
1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits (G, R, and X) enter into the computation.
3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}

\section*{Version 3.0}
\begin{tabular}{l}
\hline VSR Data Layout for xsmsub(alm)dp \\
src1 = VSR[XA] \\
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline src2 = xsmsubadp ? VSR[XT] : VSR[XB] \\
\hline DP & unused \\
\hline src3 = xsmsubadp ? VSR[XB] : VSR[XT] \\
\hline DP & unused \\
\hline tgt = VSR[XT] \\
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|c|}{ DP } \\
\hline 0 & undefined \\
\hline
\end{tabular}
\end{tabular} \begin{tabular}{l}
127 \\
\hline
\end{tabular}
\end{tabular}

\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src3} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow-\) Infinity & \(p \leftarrow-\) Infinity & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow M(\) src1,src3 \()\) & \(p \leftarrow+\) Zero & \(\mathrm{p} \leftarrow\)-Zero & \(p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow\) src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow+\) Zero & \(\mathrm{p} \leftarrow+\) Zero & \(p \leftarrow-\) Zero & \(p \leftarrow-\) Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow-\) Zero & \(\mathrm{p} \leftarrow\)-Zero & \(p \leftarrow+\) Zero & \(p \leftarrow+\) Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow\) src3 & \[
\begin{aligned}
& p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow M(\) src1,src3 \()\) & \(p \leftarrow-\) Zero & \(p \leftarrow+\) Zero & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \(p \leftarrow\) src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(\mathrm{p} \leftarrow \mathrm{SrC1}\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(\mathrm{p} \leftarrow \mathrm{src} 1\) & \[
\begin{aligned}
& p \leftarrow \text { src1 } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& p \leftarrow Q(\text { src } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& P \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& P \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q \text { (src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}

\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src2} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \[
\begin{array}{|l|}
\hline \begin{array}{l}
\mathrm{v} \leftarrow \mathrm{~d} Q N a N \\
\mathrm{vxisif} \text { flag } \leftarrow 1
\end{array} \\
\hline
\end{array}
\] & \(\mathrm{V} \leftarrow-\)-nfinity & \(\mathrm{V} \leqslant-\) - lninity & \(\mathrm{V} \leqslant-\) - lnfinity & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow-\) Infinity & v ¢ Src2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 } 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{v} \leftarrow+\) lnfinity & \(v \leftarrow S(p, s \mathrm{sc} 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow S(p, s \mathrm{sc} 2)\) & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leqslant\) Src2 & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow Q(\text { src } 2) \\
& \mathrm{vxsnn} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{v} \leftarrow+\) lnfinity & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(v \leftarrow\) Rezd & \(v \leftarrow-\) Zero & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(\mathrm{v} \leftarrow-\) Infinity & v \& SrC2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src } 2) \\
& \mathrm{vxsnn} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow+\) lnfinity & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(v \leftarrow+\) Zero & \(v \leftarrow\) Rezd & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(\mathrm{V} \leftarrow-\) Infinity & v \& SrC2 & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src2 }) \\
& v x s n a n \_f l a g \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{V} \leftarrow+\) lnfinity & \(v \leftarrow S(p, s \mathrm{sc} 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow S(p, s \mathrm{sr} 2)\) & \(\mathrm{V} \leftarrow-\) Infinity & V \& Src2 & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src } 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow+\) lnfinity & \(v \leftarrow+\) lnfinity & \(v \leftarrow+\) lnfinity & \(v \leftarrow+\) lnfinity & \(\mathrm{V} \leftarrow+\) lnfinity & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{~d} \text { NaN } \\
& \text { vxisi_flag } \leftarrow 1
\end{aligned}
\] & v \& SrC2 & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow Q(\text { srč2) } \\
& \mathrm{vxsnnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \begin{tabular}{l}
\(v \leftarrow p\) \\
vxsnan_flag \(\leftarrow 1\)
\end{tabular} \\
\hline \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & V \& SrC2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 } 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline Explanation: & \\
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
\hline src2 & For \(\boldsymbol{x s m s u b a d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT]. For \(\boldsymbol{x s m s u b m d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XB]. \\
\hline src3 & For \(\boldsymbol{x s m s u b a d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XB]. For \(\boldsymbol{x s m s u b m d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT]. \\
\hline dQNaN & Default quiet NaN (0x7FF8_0000_0000_0000). \\
\hline NZF & Nonzero finite number. \\
\hline Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. \\
\hline \(Q(x)\) & Return a QNaN with the payload of x . \\
\hline \(\mathrm{S}(\mathrm{x}, \mathrm{y})\) & Return the normalized sum of floating-point value \(x\) and negated floating-point value \(y\), having unbounded range and precision. Note: If \(x=y, v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline \(\mathrm{M}(\mathrm{x}, \mathrm{y})\) & Return the normalized product of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. \\
\hline p & The intermediate product having unbounded range and precision. \\
\hline v & The intermediate result having unbounded range and precision. \\
\hline
\end{tabular}

Table 86.Actions for xsmsub(alm)dp

VSX Scalar Multiply-Subtract Single-Precision XX3-form

xsmsubmsp \(\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 60 & & T & & A & & B \\
\hline 0 & & 6 & & 11 & & 16
\end{tabular}
```

reset_xflags()
if "xsmsubasp" then do
src1 \leftarrowVSR[32xAX+A].dword[0]
src2 }\leftarrow\operatorname{VSR[32xTX+T] .dword[0]
src3 \leftarrow VSR[32xBX+B].dword[0]
end
if "xsmsubmsp" then do
src1 \leftarrowVSR[32xAX+A].dword[0]
src2 }\leftarrowv\operatorname{VSR[32xBX+B].dword[0]
src3 \leftarrowVSR[32xTX+T].dword[0]
end
v \& MultiplyAddDP (src1,src3,NegateDP (src2))
result }\leftarrow\mathrm{ RoundToSP(RN,v)
if (vxsnan_flag) then SetFX(vXSNAN)
if(vximz_flag) then SetFX(vximZ)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(0X)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
vex_flag \leftarrow VE \& (vxsnan_flag | vximz_flag | vxisi_flag)
if( ~vex_flag ) then do
VSR[32xTX+T] .dword[0] \leftarrow ConvertSPtoSP64 (result)
VSR[32xTX+T] .dword[1] \& OxUUUU_UUUU_UUOU_UUU
FPRF }\leftarrow\mathrm{ ClassSP(result)
FR \leftarrow inc_flag
FI \&xx_flag
end
else do
FR }\leftarrow0.0
FI }\leftarrow0\textrm{bO
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

For xsmsubasp, do the following.
- Let src1 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X A]\).
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
- Let src3 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X B]\).

\section*{For xsmsubmsp, do the following.}
- Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
- Let src3 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
src1 is multiplied \({ }^{[1]}\) by src3, producing a product having unbounded range and precision.

See part 1 of Table 87, "Actions for \(x\) smsub(a|m)sp".
\(\operatorname{src} 2\) is negated and added \({ }^{[2]}\) to the product, producing a sum having unbounded range and precision.

The result, having unbounded range and precision, is normalized \({ }^{[3]}\).

See part 2 of Table 87, "Actions for \(x\) smsub(a|m)sp".
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\footnotetext{
1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( \(\mathrm{G}, \mathrm{R}\), and X ) enter into the computation.
3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}

Special Registers Altered
FPRF FR FI FX OX UX XX
VXSNAN VXISI VXIMZ

VSR Data Layout for xsmsub(alm)sp
src1 = VSR[XA]

\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{\begin{tabular}{l}
Part 1: \\
Multiply
\end{tabular}} & \multicolumn{8}{|c|}{src3} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow-\) Infinity & \(p \leftarrow-\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \mathrm{Q} \text { (src3) } \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline -NZF & \(p \leftarrow+\) Infinity & \(\mathrm{p} \leftarrow \mathrm{M}(\mathrm{src} 1, \mathrm{Src} 3)\) & \(p \leftarrow+\) Zero & \(p \leftarrow-\) Zero & \(p \leftarrow M(\) src1, src 3\()\) & \(p \leftarrow+\) Infinity & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \mathrm{Q}(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline -Zero & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{p} \leftarrow+\) Zero & \(p \leftarrow+\) Zero & \(p \leftarrow-\) Zero & \(\mathrm{p} \leftarrow-\) Zero & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline -̇. +Zero & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{dQNaN} \\
& \text { vximz__lag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{p} \leftarrow-\) Zero & \(p \leftarrow-\) Zero & \(p \leftarrow+\) Zero & \(\mathrm{p} \leftarrow+\) Zero & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{dQNaN} \\
& \text { vximz__lag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow \operatorname{Src} 3\) & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \mathrm{Q}(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline ¢ +NZF & \(p \leftarrow-\) Infinity & \(\mathrm{p} \leftarrow \mathrm{M}(\mathrm{src} 1, \mathrm{Src} 3)\) & \(p \leftarrow-\) Zero & \(p \leftarrow+\) Zero & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{array}{|l|}
\hline \mathrm{p} \leftarrow \mathrm{Q} \text { (src3) } \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline +Infinity & \(p \leftarrow-\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \mathrm{Q}(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline QNaN & \(\mathrm{p} \leftarrow \mathrm{SrCl}\) & \(\mathrm{p} \leftarrow \mathrm{SrCl}\) & \(p \leftarrow \operatorname{SrC1}\) & \(p \leftarrow \operatorname{srcl}\) & \(p \leftarrow \operatorname{SrC1}\) & \(p \leftarrow \operatorname{SrC1}\) & \(p \leftarrow \operatorname{src1}\) & \[
\begin{array}{|l|}
\hline \mathrm{p} \leftarrow \operatorname{src1} \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline SNaN & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \mathrm{Q}(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{Q}(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \mathrm{Q}(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \mathrm{Q} \text { (src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{Q}(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \mathrm{Q}(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}


\section*{Explanation:}
\begin{tabular}{|c|c|}
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
\hline src2 & For \(\boldsymbol{x s m s u b a s p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT]. For \(\boldsymbol{x s m s u b m s p}\), the double-precision floating-point value in doubleword element 0 of VSR[XB]. \\
\hline src3 & For xsmsubasp, the double-precision floating-point value in doubleword element 0 of VSR[XB]. For \(\boldsymbol{x s m s u b m s p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT]. \\
\hline dQNaN & Default quiet NaN (0x7FF8_0000_0000_0000). \\
\hline NZF & Nonzero finite number. \\
\hline Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. \\
\hline \(Q(x)\) & Return a QNaN with the payload of x . \\
\hline S(x,y) & Return the normalized sum of floating-point value \(x\) and negated floating-point value \(y\), having unbounded range and precision. Note: If \(x=y, v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline M(x,y) & Return the normalized product of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. \\
\hline p & The intermediate product having unbounded range and precision. \\
\hline \(\checkmark\) & The intermediate result having unbounded range and precision. \\
\hline
\end{tabular}

Table 87.Actions for xsmsub(a|m)sp

VSX Scalar Multiply-Subtract Quad-Precision [using round to Odd] X-form
\begin{tabular}{lll} 
xsmsubqp & VRT,VRA,VRB & \((R 0=0)\) \\
xsmsubqpo & VRT,VRA,VRB & \((R 0=1)\)
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 63 & \({ }_{6}\) VRT & VRA & VRB & & 420 & \begin{tabular}{l} 
R0 \\
31
\end{tabular} \\
\hline
\end{tabular}
```

if MSR. VSX=O then VSX_Unavailable|)
reset_xflags()
srcl \& bfp_CONVERT_FROM_ BFP128(VSR[ VRA+32])
sic2 \&bf_CONVERT_FROM_BFP128(VSR[VRT+32])
sic3 \& bfp_CONvert_froM_-bfp128(vgr[ Vrb+32])
v \&bfp_MULTIPLY_ADOIsrcl, src3, bfp_NEGATE(src2))
fnd \&bfp_ROUNDTO_BPP128(ro, FPSCR.RN, v)
result \& bfp_CONVERT_TO_BFP128(rnd)
if(vxsnan_flag) then SetFX(FPSCR.vXSNaN)
if(uximz_flag) then SetFX(FPSCR.VXIMZ)
if(ox_flag) then SetFx(FPSCR.OX)
if(ux_flag) then SetFX(FPSCR.UX)
if(xx_flag) then SetFX(FPSCR.XX)
vx_flag \leftarrowvxsnan_flag| vxim_flag
ex_flag \&FPSCR.VE\&vx_flag
if ex_flag=0 then do
vsR[[VRT+32] \& result
FPSCR.FPRF \& fprf_CLASS_BPP128(result)
end
FPSCR.FR\&(vx_flag=O) \& inc_flag
FPSCR.FI\leftarrow(vx_f |ag=0)\& <x_flag

```

Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let src2 be the floating-point value in VSR[VRT+32] represented in quad-precision format.

Let src3 be the floating-point value in VSR[VRB+32] represented in quad-precision format.

If either \(\operatorname{srcl}, \operatorname{src} 2\), or \(\operatorname{srcs}\) is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN is set to 1.

If \(\operatorname{srcl}\) is an Infinity value and \(\operatorname{srcs}\) is a Zero value, or if srcl is a Zero value and srcs is an Infinity value, an Invalid Operation exception occurs and VXI MZ is set to 1.

If \(\mathrm{srCl}_{2}\) and the product of \(\mathrm{srCl}^{2}\) and \(\mathrm{srCl}^{\text {are Infinity }}\) values having same signs, an Invalid Operation exception occurs and VXISI is set to 1 .

If srCl is a Signalling NaN , the result is the Quiet NaN corresponding to sicl.

Otherwise, if srcl is a Quiet NaN , the result is srcl .
Otherwise, if \(\mathrm{SrC2}\) is a Signalling NaN , the result is the Quiet NaN corresponding to src2.

Otherwise, if src is a Quiet NaN , the result is src .
Otherwise, if \(\mathrm{SrCl}^{\mathrm{Cl}}\) is a Signalling NaN , the result is the Quiet NaN corresponding to src .

Otherwise, if \(\mathrm{src}_{3}\) is a Quiet NaN , the result is \(\mathrm{src3}\).
Otherwise, if srCl is an Infinity value and \(\mathrm{src3}\) is a Zero value, or if \(\operatorname{srcl}\) is a Zero value and \(\operatorname{srcs}\) is an Infinity value, the result is the default Quiet \(\mathrm{NaN}^{[1]}\).

Otherwise, if the product of srcl and \(\mathrm{src3}\), and srcl are Infinity values having same signs, the result is the default Quiet NaN .

Otherwise, do the following.
srcl is multiplied by src3, producing a product having unbounded significand precision and exponent range.

See part 1 of Table 88. "Actions for xsmsubqp[o]".
srct is negated and added to the product, producing a sum having unbounded range and precision.

See part 2 of Table 88. "Actions for xsmsubqp[o]".
If the intermediate result is Tiny (i.e., the unbiased exponent is less than -16382) and \(U E=0\), the significand is shifted right \(N\) bits, where \(N\) is the difference between - 16382 and the unbiased exponent of the intermediate result. The exponent of the intermediate result is set to the value - 16382.

If \(R O=1\), let the rounding mode be Round to Odd. Otherwise, let the rounding mode be specified by RN. Unless the result is an Infinity or a Zero, the intermediate result is rounded to quad-precision using the specified rounding mode.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into VSR[ VRT +32 ] in quad-precision format.

FPRF is set to the class and sign of the result. FR is set to indicate if the rounded result was incremented. FI is set to indicate the result is inexact.

If a trap-disabled Invalid Operation exception occurs, \(F R\) and \(F I\) are set to 0 .

\section*{Version 3.0}

\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{\begin{tabular}{l}
Part 1: \\
Multiply
\end{tabular}} & \multicolumn{8}{|c|}{src3} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \multirow[t]{2}{*}{\(p \leqslant t\) nfinity} & & \multicolumn{2}{|c|}{\[
\begin{gathered}
\rho \leftarrow \text { iONIaN } \\
\text { vxim_ }
\end{gathered}
\]} & & \multirow[t]{2}{*}{\(p \leftarrow \cdot \mid n f i n i t y\)} & \multirow{6}{*}{\(p+5103\)} & \multirow{6}{*}{\[
\begin{aligned}
& p \leftarrow \text { quiet }(\operatorname{ssc} 3) \\
& \text { vxsinan_flag }+1
\end{aligned}
\]} \\
\hline -NZF & & \[
\underset{m u l}{p+\operatorname{sic} 1, \operatorname{src} 3)}
\] & \multirow[b]{2}{*}{\(p+\) Lero} & & \[
\underset{m u l(\operatorname{src} 1, \operatorname{src} 3)}{p}
\] & & & \\
\hline -Zero & \multirow[b]{2}{*}{} & & & \multicolumn{2}{|l|}{\(p \leftarrow\) - 2 ero} & \multirow[b]{2}{*}{\[
\underset{\text { vxim_ }}{\rho+\text { IONaN }}
\]} & & \\
\hline +Zero & & & \(p+\) Zero & \multicolumn{2}{|l|}{\(p+\) +2ero} & & & \\
\hline - +NZF & \multicolumn{2}{|l|}{\multirow[t]{2}{*}{\[
\begin{array}{ll}
p \leftarrow \cdot \mid \text { nfinity } & \begin{array}{l}
p \leftarrow \\
m u l|s i c 1, s r c 3| \\
\hline
\end{array} \\
\hline
\end{array}
\]}} & \multicolumn{2}{|l|}{} & \[
\underset{\operatorname{mul} \mid \operatorname{sic} 1, \operatorname{src} 3)}{ }
\] & \multirow[b]{2}{*}{\(p<t\) nfinity} & & \\
\hline +Infinity & & & & & & & & \\
\hline QNaN & \multicolumn{7}{|c|}{\(p+\operatorname{sicl}\)} & \[
\begin{gathered}
p+\text { sicl } \\
\text { vxsinan_tiag } \leqslant 1
\end{gathered}
\] \\
\hline SNaN & \multicolumn{8}{|c|}{\(p\) \&quiet (scicl) vxsnan_fag \(\leqslant 1\)} \\
\hline \multirow[t]{2}{*}{\begin{tabular}{l}
Part 2: \\
Subtract
\end{tabular}} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \[
\begin{aligned}
& v \leftarrow \mathrm{dQNaN}^{2} \\
& \text { vxisi_flag }
\end{aligned}
\] & & & & & \multirow[t]{5}{*}{\(\mathrm{V} \leftarrow-\)-nfinity} & \multirow{6}{*}{\(v \leftarrow \operatorname{src} 2\)} & \multirow{6}{*}{\[
\begin{aligned}
& V \leftarrow \text { quiet(src2) } \\
& \text { vxsnan_liag } \leftarrow 1
\end{aligned}
\]} \\
\hline -NZF & \multicolumn{2}{|r|}{\multirow[t]{3}{*}{\(V \leftarrow \operatorname{sub}(\mathrm{p}, \mathrm{Src} 2)\)

\(\mathrm{V} \leftarrow-\mathrm{Src} 2\)}} & \multicolumn{2}{|c|}{\(v \leftarrow p\)} & \(v \leftarrow \operatorname{sub}(\mathrm{p}, \mathrm{src} 2)\) & & & \\
\hline -Zero & & & \(v \leftarrow\) Rezd & \(V \leftarrow\)-Zero & \multirow{2}{*}{\(\mathrm{V} \leftarrow-\mathrm{SrC2}\)} & & & \\
\hline +Zero & & & \(v \leftarrow+\) Zero & \(v \leftarrow\) Rezd & & & & \\
\hline +NZF & & \(v \leftarrow \operatorname{sub}(\mathrm{p}, \mathrm{src} 2)\) & \multicolumn{2}{|c|}{\(V \leftarrow p\)} & \(v \leftarrow \operatorname{sub}(\mathrm{p}, \mathrm{src} 2)\) & & & \\
\hline +Infinity & \(\mathrm{V} \leftarrow+\) lnfinity & & & & & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{~d} \text { QNaN } \\
& \text { vxisi_flag } \leftarrow 1
\end{aligned}
\] & & \\
\hline QNaN \& src1 is a NaN & \multicolumn{6}{|c|}{\multirow[b]{2}{*}{\(v \leftarrow p\)}} & & \[
\begin{gathered}
V \leftarrow p \\
\text { vxsnan_flag } \leftarrow 1
\end{gathered}
\] \\
\hline QNaN \& src1 not a NaN & & & & & & & \(V \leqslant \operatorname{SrC2}\) & \[
v \leqslant \text { quiet(src2) }
\]
\[
\text { vxsnan_flag } \leftarrow 1
\] \\
\hline \multicolumn{9}{|l|}{Explanation:} \\
\hline srcl & \multicolumn{8}{|l|}{The quad-precision floating-point value in VSR[ VRA +32].} \\
\hline sic2 & \multicolumn{8}{|l|}{The quad-precision floating-point value in VSR[ VRT +32].} \\
\hline stc3 & \multicolumn{8}{|l|}{The quad-precision floating-point value in VSR[ VRB +32].} \\
\hline dQNaN & \multicolumn{8}{|l|}{} \\
\hline NZF & \multicolumn{8}{|l|}{Nonzero finite number.} \\
\hline Rezd & \multicolumn{8}{|l|}{Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands.} \\
\hline quiet ( x ) & \multicolumn{8}{|l|}{Return a QNaN with the payload of \(x\).} \\
\hline \(\operatorname{sub}(x, y)\) & \multicolumn{8}{|l|}{Return the normalized sum of floating-point value \(x\) and negated floating-point value \(y\), having unbounded range and precision. Note: If \(x=y, v\) is considered to be an exact-zero-difference result (Rezd).} \\
\hline mul ( \(x, y\) ) & \multicolumn{8}{|l|}{Return the normalized product of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision.} \\
\hline \[
p
\] & \multicolumn{8}{|l|}{The intermediate product having unbounded range and precision.} \\
\hline v & \multicolumn{8}{|l|}{The intermediate result having unbounded range and precision.} \\
\hline
\end{tabular}

\section*{| Table 88.Actions for xsmsubqp[o]}

VSX Scalar Multiply Double-Precision XX3-form


Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 x A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

Let src2 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X B]\).
src1 is multiplied \({ }^{[1]}\) by src2, producing a product having unbounded range and precision.

The product is normalized \({ }^{[2]}\).

\section*{See Table 89.}

The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

Special Registers Altered
FPRF FR FI FX OX UX XX
VXSNAN VXIMZ

\section*{VSR Data Layout for xsmuldp}
src1 = VSR[XA]
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline src2 \(=\mathrm{VSR}[\mathrm{XB}]\) \\
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)


\footnotetext{
1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{}} & \multicolumn{8}{|c|}{src2} \\
\hline & & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \multirow{8}{*}{-} & -Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(V \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline & -NZF & \(\mathrm{V} \leftarrow+\) Infinity & \(v \leftarrow M(\operatorname{src} 1, \mathrm{src} 2)\) & \(v \leftarrow+\) Zero & \(\mathrm{v} \leftarrow\)-Zero & \(v \leftarrow M(\operatorname{src} 1, \mathrm{src} 2)\) & \(V \leftarrow+\) Infinity & \(\mathrm{v} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline & -Zero & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(v \leftarrow+\) Zero & \(v \leftarrow+\) Zero & \(\mathrm{v} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow\)-Zero & \[
\begin{array}{|l|l}
\hline v \leftarrow d Q N a N \\
\text { vximz_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l}
\hline v \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline & +Zero & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(v \leftarrow\)-Zero & \(\mathrm{v} \leftarrow\)-Zero & \(v \leftarrow+\) Zero & \(v \leftarrow+\) Zero & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline & +NZF & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{M}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{v} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow+\) Zero & \(v \leftarrow M(\) src1, src2 \()\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline & +Infinity & \(\mathrm{V} \leftarrow-\) Infinity & \(V \leftarrow+\) Infinity & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline v \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2)
\]
\[
\text { vxsnan_flag } \leftarrow 1
\] \\
\hline & QNaN & \(\mathrm{v} \leftarrow \mathrm{SrCl} 1\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{v} \leftarrow \mathrm{SrC1}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{v} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \[
\begin{aligned}
& v \leftarrow \text { src1 } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline & SNaN & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src1 }) \\
& v x \text { nnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src1 }) \\
& v x \text { nnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow Q(\text { src1 }) \\
& \mathrm{vxsnan} \mathrm{\_flag} \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src1 }) \\
& v x \text { snan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow Q(\text { src } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \multicolumn{10}{|l|}{Explanation:} \\
\hline & & \multicolumn{8}{|l|}{The double-precision floating-point value in doubleword element 0 of VSR[XA].} \\
\hline & & \multicolumn{8}{|l|}{The double-precision floating-point value in doubleword element 0 of VSR[XB].} \\
\hline & & \multicolumn{8}{|l|}{Default quiet NaN (0x7FF8_0000_0000_0000).} \\
\hline & & \multicolumn{8}{|l|}{Nonzero finite number.} \\
\hline & & \multicolumn{8}{|l|}{\multirow[t]{2}{*}{Return the normalized product of floating-point value x and floating-point value y , having unbounded range and precision. Return a QNaN with the payload of x .}} \\
\hline & & & & & & & & & \\
\hline & & \multicolumn{8}{|l|}{The intermediate result having unbounded signficand precision and unbounded exponent range.} \\
\hline
\end{tabular}

Table 89.Actions for xsmuldp

VSX Scalar Multiply Quad-Precision [using
round to Odd] X-form
\begin{tabular}{lll} 
xsmulqp & VRT,VRA,VRB & \((R 0=0)\) \\
xsmulqpo & VRT,VRA,VRB & \((R 0=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \[
63
\] & \[
{ }_{6} \text { VRT }
\] & \[
{ }_{11} \text { VRA }
\] & \[
{ }_{16} \text { VRB }
\] & 21 & 36 & RO \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablell
reset \(x f 1\) ags ()
srcl \(\leftarrow\) bf p_CONVERT_FROM_BFP128(VSR[VRA+32])
srct \(\leftarrow\) bf p_CONVERT_FROM_BFP128 (VSR[VRB+32])
\(v \quad \leftarrow\) bfp_Multiplȳ (srcí, src2)
rnd \(\leftarrow\) bf R ROUND_TO BFP128(RO, FPSCR,RN, v)
result \(\leftarrow\) bf p_CONVERT_TO_BFP128(rnd)
if(vxsnan_flag) then SetFX(FPSCR.VXSNAN)
if(vximz_ \(\overline{\mathrm{f}} \mathrm{Iag})\) then SetFX(FPSCR. VXIMZ)
if(ox_flag) then SetFX(FPSCR. OX)
if(UX_-flag) then SetFX(FPSCR. UX)
if(xx_flag) then SetFX(FPSCR. XX)
vx_flag \(\leftarrow\) vxsnan_flag | vxim_flag
ex_flag \(\leftarrow\) FPSCR.VE \& Vx_flag
if ex_flag=0 then do
VŚㅠ[VRT+32] \(\leftarrow\) result
FPSCR,FPRF \(\leftarrow\) fprf_CLASS_BPP128(result)
end
FPSCR.FR \(\leftarrow\left(v x_{-} f \mid a g=0\right) \& i n c \_f \mid a g\)
FPSCR. \(F\left|\leftarrow\left(v x_{-}^{-} f \mid a g=0\right) \& x x_{-} f\right| a g\)
Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let \(\operatorname{src} 2\) be the floating-point value in VSR[VRB+32] represented in quad-precision format.

If either srcl or srcl is a Signalling NaN, an Invalid Operation exception occurs and VXSNAN is set to 1 .

If \(s r c l_{\text {cl }}\) is an Infinity value and \(s r c 2\) is a Zero value, or if srcl is a Zero value and src2 is an Infinity value, an Invalid Operation exception occurs and VXI MZ is set to 1.

If srcl is a Signalling NaN , the result is the Quiet NaN corresponding to srcl .

Otherwise, if srcl is a Quiet NaN , the result is srcl .
Otherwise, if \(\mathrm{src}_{2}\) is a Signalling NaN , the result is the Quiet NaN corresponding to sic2.

Otherwise, if src 2 is a Quiet NaN , the result is src 2 .

Otherwise, if SrCl is an Infinity value and \(\mathrm{src2}\) is a Zero value, or if srcl is a Zero value and \(\mathrm{src}^{2}\) is an Infinity value, the result is the default Quiet \(\mathrm{NaN}^{[1]}\).

Otherwise, do the following.
The normalized product of srcl multiplied by srcl is produced with unbounded significand precision and exponent range.

See Table 90. "Actions for xsmulqp[o]".
If the intermediate result is Tiny (i.e., the unbiased exponent is less than -16382) and \(U E=0\), the significand is shifted right \(N\) bits, where \(N\) is the difference between \(\cdot 16382\) and the unbiased exponent of the intermediate result. The exponent of the intermediate result is set to the value - 16382.

If \(R O=1\), let the rounding mode be Round to Odd. Otherwise, let the rounding mode be specified by RN. Unless the result is an Infinity or a Zero, the intermediate result is rounded to quad-precision using the specified rounding mode.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into VSR[VRT+32] in quad-precision format.

FPRF is set to the class and sign of the result. \(F R\) is set to indicate if the rounded result was incremented. FI is set to indicate the result is inexact.

If a trap-disabled Invalid Operation exception occurs, \(F R\) and \(F I\) are set to 0 .

If a trap-enabled Invalid Operation exception occurs, VSR[ VRT +32 ] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered:}

FPRF FR FI FX VXSNAN VXIMZ OX UX XX

\footnotetext{
【 1. The quad-precision default Quiet NaN is the value, \(0 \times 7 F F F_{1} 8000 \_0000 \_0000 \_0000 \_0000 \_0000\).
}


\section*{| Table 90. Actions for xsmulqp[o]}

I 1. Floating-point multiplication is based on exponent addition and multiplication of the significands.

VSX Scalar Multiply Single-Precision XX3-form
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{xsmulsp \(\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)} \\
\hline \[
0^{60}
\] & \({ }_{6}{ }^{\text {T }}\) & & & B & 21 & 16 & \[
\begin{aligned}
& 4 \times B \times 1 T X \\
& 293031
\end{aligned}
\] \\
\hline
\end{tabular}
```

reset_xflags()
src1 \leftarrow VSR[32\timesAX+A].dword[0]
src2 }\leftarrow\operatorname{VSR[32\timesBX+B].dword[0]
v }\leftarrow\mathrm{ MultiplyDP(src1,src2)
result }\leftarrow\mathrm{ RoundToSP(RN,v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vximz_flag) then SetFX(VXIMZ)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
vex_flag \leftarrow VE \& (vxsnan_flag | vximz_flag)
if( ~vex_flag ) then do
VSR[32\timesTX+T].dword[0] \leftarrow ConvertSPtoSP64 (result)
VSR[32xTX+T].dword[1] \& 0xUUUU_UUUU_UUUU_UUUU
FPRF }\leftarrowClassSP(result
FR \leftarrow inc_flag
FI }\leftarrowxx_fla
end
else do
FR}\leftarrow0\textrm{ObO
FI}\leftarrow\textrm{ObO
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src1 is multiplied \({ }^{[1]}\) by src2, producing a product having unbounded range and precision.

The product is normalized \({ }^{[2]}\).
See Table 91, "Actions for xsmulsp," on page 607.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to indicate if the result was incremented when rounded. FI is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered}

FPRF FR FI FX OX UX XX VXSNAN VXIMZ

VSR Data Layout for xsmulsp
\(\operatorname{src} 1=\operatorname{VSR}[\mathrm{XA}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline src2 \(=\) VSR[XB] & unused \\
\hline DP &
\end{tabular}
tgt \(=\) VSR[XT]


\footnotetext{
1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \[
\begin{array}{|l|l}
\hline \mathrm{V} \leftarrow \mathrm{dQNaN} \\
\text { vxim__flag } \leftarrow 1 \\
\hline
\end{array}
\] & \[
\begin{array}{|l|}
\hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
\text { vximz_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { (src2 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline -NZF & \(V \leftarrow+\) Infinity & \(v \leftarrow M(\operatorname{src} 1, \mathrm{src} 2)\) & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow \mathrm{M}(\mathrm{src} 1, \mathrm{src} 2)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline -Zero & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow\)-Zero & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz__lag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l|}
\hline v \leftarrow Q(s r c 2) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline -̇. +Zero & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{V} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow\)-Zero & \(v \leftarrow+\) Zero & \(\mathrm{v} \leftarrow+\) Zero & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& v \leftarrow Q(\text { (src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline ¢ +NZF & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{M}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{V} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{M}(\operatorname{src} 1, \mathrm{src} 2)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l|}
\hline \mathrm{v} \leftarrow \mathrm{Q} \text { (src2) } \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline +Infinity & \(\mathrm{V} \leftarrow-\) Infinity & \(V \leftarrow+\) Infinity & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{dQNaN} \\
\text { vximz_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \(\mathrm{V} \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \begin{tabular}{l}
\[
\mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2)
\] \\
vxsnan_flag \(\leftarrow 1\)
\end{tabular} \\
\hline QNaN & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrCl}\) & \(\mathrm{V} \leftarrow \mathrm{SrC1}\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{src} 1 \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline SNaN & \begin{tabular}{l}
\[
\mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1)
\] \\
vxsnan_flag \(\leftarrow 1\)
\end{tabular} & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \[
\begin{aligned}
& \hline v \leftarrow Q(\operatorname{src} 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \begin{tabular}{l}
\[
\mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1)
\] \\
vxsnan_flag \(\leftarrow 1\)
\end{tabular} & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{array}{|l|}
\hline v \leftarrow Q(\operatorname{src} 1) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] & \[
\begin{aligned}
& v \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \multicolumn{9}{|l|}{Explanation:} \\
\hline src1 & \multicolumn{8}{|l|}{The double-precision floating-point value in doubleword element 0 of VSR[XA].} \\
\hline src2 & \multicolumn{8}{|l|}{The double-precision floating-point value in doubleword element 0 of VSR[XB].} \\
\hline dQNaN & \multicolumn{8}{|l|}{Default quiet NaN (0x7FF8_0000_0000_0000).} \\
\hline NZF & \multicolumn{8}{|l|}{Nonzero finite number.} \\
\hline M ( \(\mathrm{x}, \mathrm{y}\) ) & \multicolumn{8}{|l|}{\multirow[t]{2}{*}{Return the normalized product of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision.}} \\
\hline \(Q(x)\) & & & & & & & & \\
\hline \(v\) & \multicolumn{8}{|l|}{The intermediate result having unbounded signficand precision and unbounded exponent range.} \\
\hline
\end{tabular}

Table 91.Actions for xsmulsp

\section*{Version 3.0}

VSX Scalar Negative Absolute Double-Precision XX2-form


Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of doubleword element 0 of \(\operatorname{VSR}[X B]\), with bit 0 set to 1 , is placed into doubleword element 0 of VSR[XT].

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

\section*{Special Registers Altered \\ None}

VSR Data Layout for xsnabsdp
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\) VSR[XT]


VSX Scalar Negative Absolute Quad-Precision X-form
xsnabsqp VRT,VRB

if MSR. VSX \(=0\) then VSX_Unavailablell
VSR[VRT+32] \(\leftarrow V S R[V R B+32] \mid 0 \times 8000 \_00000_{2} 0000 \_0000_{\mathbf{Z}} 0000 \_0000 \_0000 \_0000\)
Let src be the floating-point value in VSR[VRB+32] represented in quad-precision format.

The negative absolute value of \(\operatorname{src}\) is placed into VSR[ VRT +32] in quad-precision format.

\section*{Special Registers Altered:}

None
| VSR Data Layout for xsnabsqp
| VSR[VRB+32]
I \(\square\)
I VSR[VRT+32]
I
\(\square \operatorname{tgt}\)

Programming Note
This instruction can be used to operate on a single-precision source operand.

\section*{VSX Scalar Negate Double-Precision XX2-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{9}{|l|}{xsnegdp \(\mathrm{XT,XB}\)} \\
\hline \[
\begin{array}{ll} 
& 60 \\
0 & \\
\hline
\end{array}
\] & 6 & T & \[
{ }_{11} \quad \text { III }
\] & 16 & B & 21 & 377 & \(\left\lvert\, \begin{aligned} & \text { BX } \\ & 30 \\ & 30 \times 1 \\ & 31\end{aligned}\right.\) \\
\hline \multicolumn{9}{|l|}{XT \(\quad \leftarrow \mathrm{TX} \| \mathrm{T}\)} \\
\hline \multicolumn{9}{|l|}{} \\
\hline \multicolumn{9}{|l|}{result \(\{0: 63\} \leftarrow \sim \operatorname{VSR}[\mathrm{XB}]\{0\}|\mid \operatorname{VSR}[\mathrm{XB}]\{1: 63\}\)} \\
\hline \multicolumn{9}{|l|}{VSR[XT] \(\leftarrow\) result \| 0xUUUU_UUUU_UUUU_UUUU} \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of doubleword element 0 of VSR[XB], with bit 0 complemented, is placed into doubleword element 0 of VSR[XT].

The contents of doubleword element 1 of VSR[XT] are undefined.

Special Registers Altered
None

\section*{VSR Data Layout for xsnegdp}
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\) VSR[XT]
\begin{tabular}{|l|l|}
\hline & \multicolumn{2}{c|}{ DP undefined } \\
\hline 0 & 64
\end{tabular}

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

\section*{VSX Scalar Negative Multiply-Add Double-Precision XX3-form}
\[
\text { xsnmaddadp } \quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}
\]
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline -60 & 6 & T & 11 & A & 16 & B & 21 & 161 & \(\left\lvert\, \begin{aligned} & \text { 2x| } 8 \times 7 \times \\ & 293031\end{aligned}\right.\) \\
\hline
\end{tabular}
xsnmaddmdp XT,XA,XB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & 11 & A & 16 & B & 21 & 169 & \(\mid \mathrm{AXBXTX}\) \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline XT & \(\leftarrow \mathrm{TX} \| \mathrm{T}\) \\
\hline XA & \(\leftarrow A X \| A\) \\
\hline XB & \(\leftarrow B X|\mid B\) \\
\hline \multicolumn{2}{|l|}{reset_xflags()} \\
\hline src1 & \(\leftarrow \mathrm{VSR}[\mathrm{XA}\) ] \(\{0: 63\}\) \\
\hline src2 & \(\leftarrow{ }^{\text {cxsnmaddadp" }}\) ? VSR[XT] \(\left.00: 63\right\}\) : VSR[XB] \(\{0: 63\}\) \\
\hline & \(\leftarrow\) "xsnmaddadp" ? VSR[XB] \{0:63\} : VSR[XT] \{0:63\} \\
\hline v\{0:inf\} & \(\leftarrow\) MultiplyAddDP(src1, src3,src2) \\
\hline \multicolumn{2}{|l|}{result \(\{0: 63\} \leftarrow \operatorname{NegateDP}(\) RoundTodP(RN, v) \()\)} \\
\hline \multicolumn{2}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{2}{|l|}{if(vximz_flag) then SetFX(VXIMZ)} \\
\hline \multicolumn{2}{|l|}{if(vxisi_flag) then SetFX(VXISI)} \\
\hline \multicolumn{2}{|l|}{if(ox_flag) then SetFX(0x)} \\
\hline \multicolumn{2}{|l|}{if(ux_flag) then SetFX(UX)} \\
\hline \multicolumn{2}{|l|}{if(xx_flag) then SetFX(XX)} \\
\hline \multicolumn{2}{|l|}{vex_flag \(\leftarrow\) VE \& (vxsnan_flag | vximz_flag | vxisi_flag)} \\
\hline \multicolumn{2}{|l|}{if( ~vex_flag ) then do} \\
\hline \multicolumn{2}{|l|}{VSR[XT] ヶ result \|| 0xUUUU_UUUU_UUUU_UUUU} \\
\hline \multicolumn{2}{|l|}{FPRF \(\leftarrow\) ClassDP(result)} \\
\hline \multicolumn{2}{|l|}{FR \(\leftarrow\) inc_flag} \\
\hline \multicolumn{2}{|l|}{FI \(\leftarrow\) xx_flag} \\
\hline \multicolumn{2}{|l|}{end} \\
\hline \multicolumn{2}{|l|}{else do} \\
\hline FR \(\leftarrow\) & \(\leftarrow 0\) \\
\hline \(\mathrm{FI} \leqslant\) & \(\leftarrow 0\) \\
\hline end & \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X A]\).

For xsnmaddadp, do the following.
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
- Let src3 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X B]\).

For xsnmaddmdp, do the following.
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
- Let src3 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X T]\).
src1 is multiplied \({ }^{[1]}\) by src3, producing a product having unbounded range and precision.

See part 1 of Table 92.
src2 is added \({ }^{[2]}\) to the product, producing a sum having unbounded range and precision.

The sum is normalized \({ }^{[3]}\).
See part 2 of Table 92.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is negated and placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 93, "Scalar Floating-Point Final Result with Negation," on page 613.

\section*{Special Registers Altered}
```

FPRF FR FI FX OX UX XX
VXSNAN VXISI VXIMZ

```

\footnotetext{
1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( \(\mathrm{G}, \mathrm{R}\), and X ) enter into the computation.
3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}

VSR Data Layout for xsnmadd(alm)dp
src1 = VSR[XA]
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
src2 \(=\) xsnmaddadp \(? \mathrm{VSR}[\mathrm{XT}]: \mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline src3 = xsnmaddadp ? VSR[XB] : VSR[XT] \\
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline DP & \multicolumn{2}{c|}{ undefined } \\
\hline 0 & 64
\end{tabular}

\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src3} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow-\) Infinity & \(p \leftarrow-\) Infinity & \(p \leftarrow\) src3 & \[
\begin{aligned}
& p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow M(s r c 1, s r c 3)\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow\) src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow+\) Zero & \(\mathrm{p} \leftarrow+\) Zero & \(\mathrm{p} \leftarrow-\) Zero & \(p \leftarrow-\) Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{p} \leftarrow\)-Zero & \(\mathrm{p} \leftarrow\)-Zero & \(\mathrm{p} \leftarrow+\) Zero & \(p \leftarrow+\) Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow M(s r c 1, s r c 3)\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow\) src3 & \[
\begin{aligned}
& \mathrm{p} \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \mathrm{p} \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \[
\begin{aligned}
& p \leftarrow \text { src1 } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{P} \leftarrow Q(\text { src } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{p} \leftarrow Q(\text { src } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{p} \leftarrow Q(\text { src } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}

\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src2} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \[
\begin{array}{|l|}
\hline v \leftarrow d Q N a N \\
\text { vxisi_flag } \leftarrow 1
\end{array}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(v \leftarrow A(p, s r c 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow A\left(p, s c^{\prime} 2\right)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{v} \leftarrow \mathrm{SrC2}\) & \(\mathrm{V} \leftarrow\)-Zero & \(v \leftarrow\) Rezd & \(\mathrm{V} \leftarrow \mathrm{src} 2\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{src} 2\) & \[
\begin{array}{|l|}
\hline v \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \(v \leftarrow\) Rezd & \(\mathrm{v} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{src} 2\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(v \leftarrow A(p, s r c 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow A(p, s r c 2)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{dQNaN} \\
& \text { vxisi_flag } \leftarrow 1
\end{aligned}
\] & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \[
\begin{aligned}
& v \leftarrow p \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}

\section*{Explanation:}
\begin{tabular}{|c|c|}
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
\hline src2 & For \(\boldsymbol{x s n m a d d a d p , ~ t h e ~ d o u b l e - p r e c i s i o n ~ f l o a t i n g - p o i n t ~ v a l u e ~ i n ~ d o u b l e w o r d ~ e l e m e n t ~} 0\) of VSR[XT]. For \(\boldsymbol{x s n m a d d m d p , ~ t h e ~ d o u b l e - p r e c i s i o n ~ f l o a t i n g - p o i n t ~ v a l u e ~ i n ~ d o u b l e w o r d ~ e l e m e n t ~} 0\) of VSR[XB]. \\
\hline src3 & For \(\boldsymbol{x s n m a d d a d p , ~ t h e ~ d o u b l e - p r e c i s i o n ~ f l o a t i n g - p o i n t ~ v a l u e ~ i n ~ d o u b l e w o r d ~ e l e m e n t ~} 0\) of VSR[XB]. For xsnmaddmdp, the double-precision floating-point value in doubleword element 0 of VSR[XT]. \\
\hline dQNaN & Default quiet NaN (0x7FF8_0000_0000_0000). \\
\hline NZF & Nonzero finite number. \\
\hline Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. \\
\hline \(Q(x)\) & Return a QNaN with the payload of x . \\
\hline A(x,y) & Return the normalized sum of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. Note: If \(x=-y, v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline M ( \(\mathrm{x}, \mathrm{y}\) ) & Return the normalized product of floating-point value \(x\) and floating-point value y , having unbounded range and precision. \\
\hline p & The intermediate product having unbounded range and precision. \\
\hline \(v\) & The intermediate result having unbounded range and precision. \\
\hline
\end{tabular}

Table 92.Actions for xsnmadd(alm)dp
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline Case & ш & Ш & แ & 亗 & \[
\underset{\times}{\boldsymbol{x}}
\] &  &  &  &  &  &  &  & Returned Results and Status Setting \\
\hline \multirow{9}{*}{Special} & - & - & - & - & - & 0 & 0 & 0 & - & - & - & - & T(N(r)), FPRF \(<\) ClassFP(r), F1<0, FR \(\leftarrow 0\) \\
\hline & 0 & - & - & - & - & - & - & 1 & - & - & - & - & T( \(r\) ), FPRF \(\leftarrow\) ClassFP( \((\mathrm{r}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{fx}(\mathrm{VXISI})\) \\
\hline & 0 & - & - & - & - & 0 & 1 & - & - & - & - & - &  \\
\hline & 0 & - & - & - & - & 1 & 0 & - & - & - & - & - & \(\mathrm{T}(\mathrm{r}), \mathrm{FPRF} \leftarrow C \mathrm{ClassFP}(\mathrm{r}), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0, \mathrm{fx}(\mathrm{VXSNAN})\) \\
\hline & 0 & - & - & - & - & 1 & 1 & - & - & - & - & - & \(\mathrm{T}(\mathrm{r}), \mathrm{FPRF} \leftarrow<\) ClassFP(r), F1<0, FR \(<0, \mathrm{fx}(\) VXSNAN), fx(VXIMZ) \\
\hline & 1 & - & - & - & - & - & - & 1 & - & - & - & - & fx(VXISII), error() \\
\hline & 1 & - & - & - & - & 0 & 1 & - & - & - & - & - & fx(VXIMZ), error() \\
\hline & 1 & - & - & - & - & 1 & 0 & - & - & - & - & - & fx(VXSNAN), error() \\
\hline & 1 & - & - & - & - & 1 & 1 & - & - & - & - & - & fx(VXSNAN), fx(VXIMZ), error() \\
\hline \multirow{5}{*}{Normal} & - & - & - & - & - & - & - & - & no & - & - & - & T(N(r)), FPRF¢ClassFP \((N(r)), \mathrm{Fl} \leftarrow 0, \mathrm{FR} \leftarrow 0\) \\
\hline & - & - & - & - & 0 & - & - & - & yes & no & - & - & T \((N(r)\) ), FPRF \(\leftarrow C \operatorname{ClassFP}(N(r)\) ), Fl \(<1, \mathrm{FR} \leftarrow 0, \mathrm{fx}(\mathrm{XX})\) \\
\hline & - & - & - & - & 0 & - & - & - & yes & yes & - & - &  \\
\hline & - & - & - & - & 1 & - & - & - & yes & no & - & - & T(N( \((\mathrm{r})\) ), FPRF \(<\) ClassFP( \((\mathrm{N}(\mathrm{r})\), Fl< \(<1, \mathrm{FR} \leftarrow 0, \mathrm{fx}(\mathrm{XX})\), error() \\
\hline & - & - & - & - & 1 & - & - & - & yes & yes & - & - & T(N( \((\mathrm{r})\) ), FPRF \(<\) ClassFP( \((\mathrm{N}(\mathrm{r})\), \(\mathrm{F} /<1, \mathrm{FR} \leftarrow 1, \mathrm{fx}(\mathrm{XX})\), error() \\
\hline \multirow{5}{*}{Overflow} & - & 0 & - & - & 0 & - & - & - & - & - & - & - &  \\
\hline & - & 0 & - & - & 1 & - & - & - & - & - & - & - &  \\
\hline & - & 1 & - & - & - & - & - & - & - & - & no & - &  \\
\hline & - & 1 & - & - & - & - & - & - & - & - & yes & no &  \\
\hline & - & 1 & - & - & - & - & - & - & - & - & yes & yes &  \\
\hline
\end{tabular}

Explanation:
\begin{tabular}{|c|c|}
\hline - & The results do not depend on this condition. \\
\hline ClassFP(x) & Classifies the floating-point value x as defined in Table 2, "Floating-Point Result Flags," on page 373. \\
\hline \(\mathrm{fx}(\mathrm{x})\) & FX is set to 1 if \(\mathrm{x}=0 . \mathrm{x}\) is set to 1 . \\
\hline \(\beta\) & Wrap adjust, where \(\beta=2^{1536}\) for double-precision and \(\beta=2^{192}\) for single-precision. \\
\hline q & The value defined in Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516, signficand rounded to the target precision, unbounded exponent range. \\
\hline \(r\) & The value defined in Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516, signficand rounded to the target precision, bounded exponent range. \\
\hline \(v\) & The precise intermediate result defined in the instruction having unbounded signficand precision, unbounded exponent range. \\
\hline FI & Floating-Point Fraction Inexact status flag, FPSCR \({ }_{\text {Fl }}\). This status flag is nonsticky. \\
\hline FR & Floating-Point Fraction Rounded status flag, \(\mathrm{FPSCR}_{\text {FR }}\). \\
\hline OX & Floating-Point Overflow Exception status flag, FPSCR \({ }_{\text {Ox }}\). \\
\hline error() & The system error handler is invoked for the trap-enabled exception if the FE0 and FE1 bits in the Machine State Register are set to any mode other than the ignore-exception mode. \\
\hline \(N(x)\) & The value x is is negated by complementing the sign bit of x . \\
\hline T(x) & The value x is placed in element 0 of VSR[XT] in the target precision format. The contents of the remaining element(s) of VSR[XT] are undefined. \\
\hline UX & Floating-Point Underflow Exception status flag, FPSCR \({ }_{\text {Ux }}\) \\
\hline VXSNAN & Floating-Point Invalid Operation Exception (SNaN) status flag, FPSCR \({ }_{\text {VxSNAN }}\). \\
\hline VXIMZ & Floating-Point Invalid Operation Exception (Infinity \(\times\) Zero) status flag, FPSCR \({ }_{\text {VxImz }}\). \\
\hline VXISI & Floating-Point Invalid Operation Exception (Infinity - Infinity) status flag, FPSCR \({ }_{\text {VxISI }}\) - \\
\hline XX & Float-Point Inexact Exception status flag, FPSCR \({ }_{X x}\). The flag is a sticky version of FPSCR \(_{\text {FI }}\). When FPSCR \(_{F I}\) is set to a new value, the new value of FPSCR \(_{X X}\) is set to the result of ORing the old value of FPSCR \(_{X X}\) with the new value of \(F P S C R_{\text {FI }}\). \\
\hline
\end{tabular}

Table 93.Scalar Floating-Point Final Result with Negation

\section*{Version 3.0}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline Case & ш & Ш & 山 & N & \[
\underset{\times}{\boldsymbol{x}}
\] &  &  &  &  & &  &  &  & Returned Results and Status Setting \\
\hline \multirow{8}{*}{Tiny} & - & - & 0 & - & - & - & - & - & no & & - & - & - & T(N(r)), FPRF \(\leftarrow\) ClassFP( \((\mathrm{r}(\mathrm{r})\), Fl<0, FR \(\leftarrow 0\) \\
\hline & - & - & 0 & - & 0 & - & - & - & yes & & no & - & - & T(N(r)), FPRF \(\leftarrow\) ClassFP( \((\mathrm{N}(\mathrm{r})\), FF<-1, FR \(\leftarrow 0, \mathrm{fx}(\mathrm{UX})\), fx(XX) \\
\hline & - & - & 0 & - & 0 & - & - & - & yes & & yes & - & - &  \\
\hline & - & - & 0 & - & 1 & - & - & - & yes & & no & - & - &  \\
\hline & - & - & 0 & - & 1 & - & - & - & yes & & yes & - & - & T(N(r)), FPRF<ClassFP( \((\) (r) ), Fl< \(<1, \mathrm{FR} \leftarrow 1, \mathrm{fx}(\mathrm{UX})\), fx(XX), error \()\) \\
\hline & - & - & 1 & - & - & - & - & - & yes & & - & no & - &  \\
\hline & - & - & 1 & - & - & - & - & - & yes & & - & yes & no &  \\
\hline & - & - & 1 & - & - & - & - & - & yes & & - & yes & yes &  \\
\hline
\end{tabular}

Explanation:
\begin{tabular}{|c|c|}
\hline - & The results do not depend on this condition. \\
\hline ClassFP(x) & Classifies the floating-point value x as defined in Table 2, "Floating-Point Result Flags," on page 373. \\
\hline fx(x) & \(F X\) is set to 1 if \(x=0 . x\) is set to 1 . \\
\hline \(\beta\) & Wrap adjust, where \(\beta=2^{1536}\) for double-precision and \(\beta=2^{192}\) for single-precision. \\
\hline q & The value defined in Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516, signficand rounded to the target precision, unbounded exponent range. \\
\hline \(r\) & The value defined in Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516, signficand rounded to the target precision, bounded exponent range. \\
\hline \(v\) & The precise intermediate result defined in the instruction having unbounded signficand precision, unbounded exponent range. \\
\hline FI & Floating-Point Fraction Inexact status flag, \(\mathrm{FPSCR}_{\mathrm{FI}}\). This status flag is nonsticky. \\
\hline FR & Floating-Point Fraction Rounded status flag, FPSCR \({ }_{\text {FR }}\). \\
\hline OX & Floating-Point Overflow Exception status flag, FPSCR \({ }_{\text {Ox }}\). \\
\hline error() & The system error handler is invoked for the trap-enabled exception if the FE0 and FE1 bits in the Machine State Register are set to any mode other than the ignore-exception mode. \\
\hline \(N(x)\) & The value x is is negated by complementing the sign bit of x . \\
\hline T(x) & The value x is placed in element 0 of VSR[XT] in the target precision format. The contents of the remaining element(s) of VSR[XT] are undefined. \\
\hline UX & Floating-Point Underflow Exception status flag, FPSCR \({ }_{\text {UX }}\) \\
\hline VXSNAN & Floating-Point Invalid Operation Exception (SNaN) status flag, FPSCR \({ }_{\text {vxsnan }}\). \\
\hline VXIMZ & Floating-Point Invalid Operation Exception (Infinity \(\times\) Zero) status flag, FPSCR \({ }_{\text {VxIMz }}\). \\
\hline VXISI & Floating-Point Invalid Operation Exception (Infinity - Infinity) status flag, FPSCR \({ }_{\text {VxISI }}\) - \\
\hline XX & Float-Point Inexact Exception status flag, FPSCR \({ }_{\text {Xxx }}\). The flag is a sticky version of \(\mathrm{FPSCR}_{\mathrm{FI}}\). When \(\mathrm{FPSCR}_{\mathrm{FI}}\) is set to a new value, the new value of FPSCR \(X_{X X}\) is set to the result of ORing the old value of FPSCR \({ }_{X X}\) with the new value of FPSCR \({ }_{F I}\). \\
\hline
\end{tabular}

Table 93.Scalar Floating-Point Final Result with Negation (Continued)

\section*{VSX Scalar Negative Multiply-Add Single-Precision XX3-form}
```

xsnmaddasp XT,XA,XB

```
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & & A & 16 & B & 21 & 129 & \begin{tabular}{l}
AxBx|Tx \\
293031
\end{tabular} \\
\hline \multicolumn{9}{|l|}{xsnmaddmsp \(\quad \mathrm{XT}, \mathrm{X}\)} \\
\hline
\end{tabular}

```

reset_xflags()
if "xsnmaddasp" then do
src1 \leftarrow VSR[32xAX+A].dword[0]
src2 \leftarrowVSR[32xTX+T].dword[0]
src3 \leftarrow VSR[32xBX+B].dword[0]
end
if "xsnmaddmsp" then do
src1 \leftarrow VSR[32xAX+A].dword[0]
src2 }\leftarrow\textrm{VSR[32\timesBX+B].dword[0]
src3 \leftarrow VSR[32xTX+T].dword[0]
end
v }\leftarrowMultiplyAddDP(SrC1,src3,src2)
result }\leftarrow\operatorname{NegateSP(RoundToSP(RN,v))
if(vxsnan_flag) then SetFX (VXSNAN)
if(vximz_flag) then SetFX(VXIMZ)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
vex_flag \leftarrow VE \& (vxsnan_flag | vximz_flag | vxisi_flag)
if( ~vex_flag ) then do
VSR[32\timesTX+T] .dword[0] \leftarrow ConvertToSP(result)
VSR[32xTX+T].dword[1] \& 0xUUUU_UUUU_UUUU_UUUU
FPRF \leftarrow ClassSP(result)
FR \leftarrow inc_flag
FI }\leftarrow\mathrm{ xx_flag
end
else do
FR}\leftarrow0\textrm{ObO
FI}\leftarrow\textrm{ObO
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

For xsnmaddasp, do the following.
- Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
- Let src3 be the double-precision floating-point value in doubleword element 0 of VSR[XB].

For xsnmaddmsp, do the following.
- Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
- Let src3 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
src1 is multiplied \({ }^{[1]}\) by src3, producing a product having unbounded range and precision.

See part 1 of Table 94, "Actions for xsnmadd(a|m)sp," on page 617.
src2 is added \({ }^{[2]}\) to the product, producing a sum having unbounded range and precision.

The sum is normalized \({ }^{[3]}\).
See part 2 of Table 94, "Actions for xsnmadd(a|m)sp," on page 617.

The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is negated and placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 93, "Scalar Floating-Point Final Result with Negation," on page 613.

\footnotetext{
1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( \(\mathrm{G}, \mathrm{R}\), and X ) enter into the computation.
3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}

\section*{Version 3.0}
```

Special Registers Altered
FPRF FR FI FX OX UX XX
VXSNAN VXISI VXIMZ

```


\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src3} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz__lag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow-\) Infinity & \(p \leftarrow-\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)\) & \(p \leftarrow \operatorname{srCl}\) & \(p \leftarrow \operatorname{src1}\) & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow\) src3 & \[
\begin{aligned}
& p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \hline p \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow+\) Zero & \(p \leftarrow+\) Zero & \(p \leftarrow-\) Zero & \(p \leftarrow-\) Zero & \[
\begin{aligned}
& \hline p \leftarrow \mathrm{dQNaN} \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{p} \leftarrow \mathrm{Src} 3\) & \[
\begin{aligned}
& \hline p \leftarrow Q(\operatorname{src} 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz__lag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow-\) Zero & \(\mathrm{p} \leftarrow-\) Zero & \(\mathrm{p} \leftarrow+\) Zero & \(p \leftarrow+\) Zero & \[
\begin{aligned}
& \mathrm{p} \leftarrow \mathrm{dQNaN} \\
& \text { vximz__lag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow\) Src3 & \[
\begin{array}{|l|}
\hline p \leftarrow Q(s r c 3) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)\) & \(p \leftarrow \operatorname{src1}\) & \(\mathrm{p} \leftarrow \mathrm{SrCl}\) & \(p \leftarrow M(\operatorname{src} 1, \mathrm{SrC} 3)\) & \(p \leftarrow+\) Infinity & \(\mathrm{p} \leftarrow \mathrm{SrC3}\) & \[
\begin{aligned}
& \hline p \leftarrow Q \text { Qsrc3) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz__lag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{array}{|l|}
\hline p \leftarrow Q(s r c 3) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{SrC1}\) & \(p \leftarrow \operatorname{SrC1}\) & \(p \leftarrow \operatorname{SrC1}\) & \(p \leftarrow \operatorname{SrC1}\) & \(p \leftarrow \operatorname{src1}\) & \(p \leftarrow \operatorname{SrC1}\) & \[
\begin{aligned}
& \hline p \leftarrow \operatorname{src} 1 \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow \mathrm{Q}(\text { src1) } \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{array}{|l}
\hline p \leftarrow Q(\text { src1 }) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] & \[
\begin{aligned}
& \hline p \leftarrow Q(\operatorname{src1}) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{array}{|l|}
\hline p \leftarrow Q(\operatorname{src} 1) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \[
\begin{array}{|l}
\hline p \leftarrow Q(\operatorname{src} 1) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{array}{|l}
\hline p \leftarrow Q(\operatorname{src} 1) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] & \[
\begin{aligned}
& \hline p \leftarrow Q(\operatorname{src} 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src2} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \(\mathrm{V} \leftarrow\) - Infinity & \(\mathrm{V} \leftarrow\)-Infinity & \(V \leftarrow\)-Infinity & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow-\) Infinity & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{dQNaN} \\
\text { vxisi_flag } \leftarrow 1
\end{array}
\] & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline \(\mathrm{V} \leftarrow-\) Infinity & \(v \leftarrow A(p, s r c 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow A(p, s r c 2)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \operatorname{src} 2\) & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{V} \leftarrow\)-Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \(\mathrm{V} \leftarrow\)-Zero & \(\mathrm{V} \leftarrow\) Rezd & \(\mathrm{V} \leftarrow \operatorname{src} 2\) & \(\mathrm{V} \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{Src} 2\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] \\
\hline \(\mathrm{V} \leftarrow\) - Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \(\mathrm{V} \leftarrow\) Rezd & \(\mathrm{V} \leftarrow+\) Zero & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \(\mathrm{V} \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{Src} 2\) & \[
\begin{array}{|l|}
\hline v \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline \(\mathrm{V} \leftarrow\) - Infinity & \(v \leftarrow A(p, s r c 2)\) & \(v \leftarrow p\) & \(V \leftarrow p\) & \(v \leftarrow A(p, s r c 2)\) & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] \\
\hline \[
\begin{array}{|l|l}
\hline \mathrm{V} \leftarrow \mathrm{dQNaN} \\
\text { vxisi_flag } \leftarrow 1 \\
\hline
\end{array}
\] & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{Src} 2\) & \[
\begin{array}{|l|}
\hline v \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline \(V \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \[
\begin{aligned}
& v \leftarrow p \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(V \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\operatorname{src} 2) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}

Explanation:
\begin{tabular}{|c|c|}
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
\hline src2 & For \(\boldsymbol{x s n m a d d a s p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT]. For \(\boldsymbol{x s n m a d d m s p , ~ t h e ~ d o u b l e - p r e c i s i o n ~ f l o a t i n g - p o i n t ~ v a l u e ~ i n ~ d o u b l e w o r d ~ e l e m e n t ~} 0\) of VSR[XB]. \\
\hline src3 & For xsnmaddasp, the double-precision floating-point value in doubleword element 0 of VSR[XB]. For xsnmaddmsp, the double-precision floating-point value in doubleword element 0 of VSR[XT]. \\
\hline dQNaN & Default quiet NaN ( \(0 \times 7 \mathrm{FF} 8\) _0000_0000_0000). \\
\hline NZF & Nonzero finite number. \\
\hline Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. \\
\hline \(Q(x)\) & Return a QNaN with the payload of x . \\
\hline A(x,y) & Return the normalized sum of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. Note: If \(x=-y\), \(v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline M ( \(\mathrm{x}, \mathrm{y}\) ) & Return the normalized product of floating-point value \(x\) and floating-point value y , having unbounded range and precision. \\
\hline p & The intermediate product having unbounded range and precision. \\
\hline v & The intermediate result having unbounded range and precision. \\
\hline
\end{tabular}

Table 94.Actions for xsnmadd(a|m)sp

VSX Scalar Negative Multiply-Add Quad-Precision [using round to Odd] X-form
\begin{tabular}{lll} 
xsnmaddqp & VRT,VRA,VRB & \((R 0=0)\) \\
xsnmaddqpo & VRT,VRA,VRB & \((R 0=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \[
63
\] & \[
6_{6} \text { VRT }
\] & \[
{ }_{11} \text { VRA }
\] & \[
{ }_{16} \text { VRB }
\] & 21 & 452 & \begin{tabular}{|r} 
RO \\
31
\end{tabular} \\
\hline
\end{tabular}
```

if MSR.VSX=0 then VSX_Unavailable(l)
reset_xflags()
srCl \leftarrowbfp_CONVERT_FROM_BFP128(VSR[VRA+32])
STC2 \&bfp_CONVERT_FROM_BFP128(VSR[VRT+32])
srC3 \&bfp_CONVERT_FROM_BFP128(VSR[VRB+32])
v \&bfp_MULTIPLY_ADD(srcl,src3,src2)
rnd <bfp_NEGATE(bfp_ROUND_TO_BFP128(RO, FPSCR, RN, v))
result \& bfp_CONVERT_TO_BFP128(rnd)
if(vxsnan_flag) then SetFX(FPSCR.VXSNAN)
if(vximz_- |ag) then SetFX(FPSCR.VXI MZ)
if(OX_flag) then SetFX(FPSCR.OX)
if(ux_flag) then SetFX(FPSCR.UX)
if(xx_flag) then SetFX(FPSCR. XX)
vx_flag \leftarrowvxsnan_flag| vxim_flag
ex_flag \leftarrowFPSCR.VE \& vx_flag
if ex_flag=0 then do
VSR[VRT+32]}\leftarrow\mathrm{ result
FPSCR.FPRF \leftarrow fprf_CLASS_BFP128(result)
end
FPSCR.FR\leftarrow(v\mp@subsup{x}{_}{\prime}f(ag=0) \& inc_flag
FPSCR.FI \leftarrow(vx_flag=0) \& xx_flag

```

Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let srcl be the floating-point value in VSR[VRT+32] represented in quad-precision format.

Let src3 be the floating-point value in VSR[VRB+32] represented in quad-precision format.

If either srcl, srcl, or srcs is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN is set to 1.

If srcl is an Infinity value and srcs is a Zero value, or if srcl is a Zero value and \(\mathrm{srcs}^{2}\) is an Infinity value, an Invalid Operation exception occurs and VXI MZ is set to 1.

If \(\mathrm{srcl}_{2}\) and the product of \(\mathrm{srcl}^{2}\) and \(\mathrm{srcs}^{\text {are Infinity }}\) values having opposite signs, an Invalid Operation exception occurs and VXISI is set to 1 .

If srcl is a Signalling NaN , the result is the Quiet NaN corresponding to srcl.

Otherwise, if \(\operatorname{src} 1\) is a Quiet NaN , the result is srcl .
Otherwise, if \(\mathrm{src}_{2}\) is a Signalling NaN , the result is the Quiet NaN corresponding to sic2.

Otherwise, if src 2 is a Quiet NaN , the result is src .
Otherwise, if \(\mathrm{srcs}^{\mathrm{rc}}\) is a Signalling NaN , the result is the Quiet NaN corresponding to src3.

Otherwise, if src 3 is a Quiet NaN , the result is src 3 .
Otherwise, if srcl is an Infinity value and srcs is a Zero value, or if srcl is a Zero value and srcs is an Infinity value, the result is the default Quiet \(\mathrm{NaN}^{[1]}\).

Otherwise, if the product of \(\mathrm{srcl}_{1}\) and srcs , and \(\mathrm{srcc}_{2}\) are Infinity values having opposite signs, the result is the default Quiet NaN.

Otherwise, do the following.
srcl is multiplied by src3, producing a product having unbounded significand precision and exponent range.

See part 1 of Table 77. "Actions for xsmadd(a|m)dp".
srce is added to the product, producing a sum having unbounded range and precision.

See part 2 of Table 77. "Actions for xsmadd(a|m)dp".

If the intermediate result is Tiny (i.e., the unbiased exponent is less than -16382) and \(U E=0\), the significand is shifted right \(N\) bits, where \(N\) is the difference between - 16382 and the unbiased exponent of the intermediate result. The exponent of the intermediate result is set to the value - 16382.

If \(R O=1\), let the rounding mode be Round to Odd. Otherwise, let the rounding mode be specified by RN. Unless the result is an Infinity or a Zero, the intermediate result is rounded to quad-precision using the specified rounding mode.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is negated and placed into VSR[VRT+32] in quad-precision format.

FPRF is set to the class and sign of the result. FR is set to indicate if the rounded result was incremented. Fl is set to indicate the result is inexact.


\section*{618}

If a trap-disabled Invalid Operation exception occurs,
\(F R\) and \(F I\) are set to 0 .
If a trap-enabled Invalid Operation exception occurs, VSR[VRT+32] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

Special Registers Altered:
FPRF FR FI
FX VXSNAN VXIMZ VXISI OX UX XX
VSR Data Layout for xsnmaddqp[o]
I VSR[VRA+32]

\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{\begin{tabular}{l}
Part 1: \\
Multiply
\end{tabular}} & \multicolumn{8}{|c|}{src3} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \multirow[t]{2}{*}{\(p \leqslant+\) ninity} & & \multicolumn{2}{|c|}{\[
\begin{gathered}
\rho+\text { doNall } \\
\text { vxim_ } \\
\text { flag }+1
\end{gathered}
\]} & & \multirow[t]{2}{*}{\(p \leftarrow \cdot \mid n+i n i t y\)} & \multirow{6}{*}{\(p+5103\)} & \multirow{6}{*}{\[
\begin{aligned}
& p \leftarrow \text { quiet }(\operatorname{sc} 33) \\
& \text { vxsnan_flag }+1
\end{aligned}
\]} \\
\hline -NZF & & \[
\stackrel{p \leftarrow}{\operatorname{Mul}(\operatorname{sic} 1, \operatorname{src} 3)}
\] & \multirow[b]{2}{*}{\(p+\) +2ero} & & \[
\stackrel{p \leftarrow}{\operatorname{Mul}(\operatorname{sic} 1, \text { sic } 3)}
\] & & & \\
\hline -Zero & \multirow[b]{2}{*}{} & & & \(p \leqslant\) Lero & & \multirow[b]{2}{*}{\[
\begin{gathered}
p \leftarrow \text { dollan } \\
\text { vxim_ } \mid \text { iag }
\end{gathered}
\]} & & \\
\hline +Zero & & & \multirow[t]{2}{*}{\(p \leftarrow\) Zero} & \(p \leqslant+\) Lero & & & & \\
\hline +NZF & \multirow[b]{2}{*}{\(p \leftarrow \cdot \mid n\) ninity} & \(p \leftarrow\) \(\operatorname{Mul}(\operatorname{sic} 1, \operatorname{sic} 3)\) & & & \(p \leftarrow\) \(\operatorname{Mul}(\operatorname{sic} 1,5 r(3)\) & \multirow[b]{2}{*}{\(p \leftarrow\) thininity} & & \\
\hline +Infinity & & & \multicolumn{2}{|c|}{\[
\begin{gathered}
\rho+\text { dONaN } \\
\text { vxim_ flag } \leqslant 1
\end{gathered}
\]} & & & & \\
\hline QNaN & \multicolumn{7}{|c|}{\(p+s i c l\)} & \[
\begin{gathered}
p \leftarrow S I C 1 \\
\text { vxsnanflag } \leqslant 1
\end{gathered}
\] \\
\hline SNaN & \multicolumn{8}{|c|}{\[
\begin{aligned}
& p<\text { quiet (scicl) } \\
& \text { vxsian_flag }+1
\end{aligned}
\]} \\
\hline
\end{tabular}


\section*{Explanation:}
\begin{tabular}{|c|c|}
\hline srcl & The quad-precision floating-point value in VSR [ VRA +32]. \\
\hline src2 & The quad-precision floating-point value in VSR [ VRT +32]. \\
\hline src3 & The quad-precision floating-point value in VSR[ VRB +32 ]. \\
\hline dQNaN & Default quiet \(\mathrm{NaN}\left(0 \times 7 F F F \_8000 \_0000 \_0000 \_0000 \_0000 \_0000\right)\). \\
\hline NZF & Nonzero finite number. \\
\hline Read & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. \\
\hline quiet ( x ) & Return a QNaN with the payload of X . \\
\hline \(\operatorname{Add}(x, y)\) & Return the normalized sum of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. Note: If \(x=\cdot y, v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline Mul ( \(\mathrm{x}, \mathrm{y}\) ) & Return the normalized product of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. \\
\hline \(p\) & The intermediate product having unbounded range and precision. \\
\hline \(v\) & The intermediate result having unbounded range and precision. \\
\hline
\end{tabular}

\section*{\| Table 95.Actions for xsnmaddqp[o]}

\section*{VSX Scalar Negative Multiply-Subtract Double-Precision XX3-form}
\[
\text { xsnmsubadp } \quad \text { XT,XA,XB }
\]

xsnmsubmdp XT,XA,XB

\begin{tabular}{|c|c|}
\hline XT & \(\leftarrow T X|\mid T\) \\
\hline XA & \(\leftarrow A X \| A\) \\
\hline XB & \(\leftarrow B X|\mid B\) \\
\hline \multicolumn{2}{|l|}{reset_xflags()} \\
\hline src1 & \(\leftarrow \mathrm{VSR}[\mathrm{XA}]\{0: 63\}\) \\
\hline src2 & \(\leftarrow \mathrm{VSR}[\mathrm{XT]}\) \{0:63\} \\
\hline src3 & \(\leftarrow \operatorname{VSR}[\mathrm{XB}\) ] \(00: 63\}\) \\
\hline src2 & \(\leftarrow\) "xsnmsubadp" ? VSR[XT]\{0:63\} : VSR[XB]\{0:63\} \\
\hline src3 & \(\leftarrow\) "xsnmsubadp" ? VSR[XB]\{0:63\} : VSR[XT]\{0:63\} \\
\hline v\{0:inf\} & \(\leftarrow\) MultiplyAddDP(src1, src3, NegateDP(src2)) \\
\hline \multicolumn{2}{|l|}{result \(\{0: 63\} \leftarrow \operatorname{NegateDP}(\) RoundTodP(RN, V\()\) )} \\
\hline \multicolumn{2}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{2}{|l|}{if(vximz_flag) then SetFX(VXIMZ)} \\
\hline \multicolumn{2}{|l|}{if(vxisi_flag) then SetFX(VXISI)} \\
\hline \multicolumn{2}{|l|}{if(ox_flag) then SetFX(0X)} \\
\hline \multicolumn{2}{|l|}{if(ux_flag) then SetFX(UX)} \\
\hline \multicolumn{2}{|l|}{if(xx_flag) then SetFX(XX)} \\
\hline \multicolumn{2}{|l|}{vex_flag \(\leftarrow\) VE \& (vxsnan_flag | vximz_flag | vxisi_flag)} \\
\hline \multicolumn{2}{|l|}{if( -vex_flag ) then do} \\
\hline \multicolumn{2}{|l|}{VSR[XT] ז result \|| 0xUUUU_UUUU_UUUU_UUUU} \\
\hline \multicolumn{2}{|l|}{FPRF \(\leftarrow\) ClassDP(result)} \\
\hline \multicolumn{2}{|l|}{FR \(\leftarrow\) inc_flag} \\
\hline \multicolumn{2}{|l|}{FI \(\leftarrow\) xx_flag} \\
\hline \multicolumn{2}{|l|}{end} \\
\hline \multicolumn{2}{|l|}{else do} \\
\hline FR \(\leftarrow\) & \(\leftarrow 0 \mathrm{bo}\) \\
\hline \(\mathrm{FI} \leftarrow\) & \(\leftarrow 0 \mathrm{bo}\) \\
\hline end & \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

For xsnmsubadp, do the following.
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
- Let src3 be the double-precision floating-point value in doubleword element 0 of VSR[XB].

For xsnmsubmdp, do the following.
- Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
- Let src3 be the double-precision floating-point value in doubleword element 0 of VSR[XT].
src1 is multiplied \({ }^{[1]}\) by src3, producing a product having unbounded range and precision.

See part 1 of Table 96.
src2 is negated and added \({ }^{[2]}\) to the product, producing a sum having unbounded range and precision.

The sum is normalized \({ }^{[3]}\).
See part 2 of Table 96.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is negated and placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 93, "Scalar Floating-Point Final Result with Negation," on page 613.

\section*{Special Registers Altered}

FPRF FR FI FX OX UX XX VXSNAN VXISI VXIMZ

\footnotetext{
1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits (G, R, and X) enter into the computation.
3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}

\section*{Version 3.0}
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{VSR Data Layout for xsnmsub(alm)dp src1 = VSR[XA]} \\
\hline DP & unused \\
\hline \multicolumn{2}{|l|}{src2 = xsnmsubadp ? VSR[XT] : VSR[XB]} \\
\hline DP & unused \\
\hline \multicolumn{2}{|l|}{src3 = xsnmsubadp ? VSR[XB] : VSR[XT]} \\
\hline DP & unused \\
\hline \multicolumn{2}{|l|}{tgt \(=\mathrm{VSR}[\mathrm{XT}]\)} \\
\hline DP & undefined \\
\hline 0 & \\
\hline
\end{tabular}

\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src3} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow-\) Infinity & \(p \leftarrow-\) Infinity & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow M(\) src1,src3 \()\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow\) src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow+\) Zero & \(\mathrm{p} \leftarrow+\) Zero & \(p \leftarrow-\) Zero & \(\mathrm{p} \leftarrow\)-Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow-\) Zero & \(\mathrm{p} \leftarrow\)-Zero & \(p \leftarrow+\) Zero & \(p \leftarrow+\) Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow\) src3 & \[
\begin{aligned}
& p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow M(\) src1,src3 \()\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& \hline p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \(p \leftarrow\) src3 & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(\mathrm{p} \leftarrow \mathrm{SrC1}\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(\mathrm{p} \leftarrow \mathrm{src} 1\) & \[
\begin{aligned}
& p \leftarrow \text { src1 } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& p \leftarrow Q(\text { src } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& P \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& P \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q \text { (src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}

\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src2} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \[
\begin{array}{|l|}
\hline \begin{array}{l}
\mathrm{v} \leftarrow \mathrm{~d} Q N a N \\
\mathrm{vxisif} \text { flag } \leftarrow 1
\end{array} \\
\hline
\end{array}
\] & \(\mathrm{V} \leftarrow-\)-nfinity & \(\mathrm{V} \leqslant-\) - lninity & \(\mathrm{V} \leqslant-\) - lnfinity & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow-\) Infinity & v ¢ Src2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 } 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{v} \leftarrow+\) lnfinity & \(v \leftarrow S(p, s \mathrm{sc} 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow S(p, s \mathrm{sc} 2)\) & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leqslant\) Src2 & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow Q(\text { src } 2) \\
& \mathrm{vxsnn} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{v} \leftarrow+\) lnfinity & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(v \leftarrow\) Rezd & \(v \leftarrow-\) Zero & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(\mathrm{v} \leftarrow-\) Infinity & v \& SrC2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src } 2) \\
& \mathrm{vxsnn} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow+\) lnfinity & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(v \leftarrow+\) Zero & \(v \leftarrow\) Rezd & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(\mathrm{V} \leftarrow-\) Infinity & v \& SrC2 & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src2 }) \\
& v x s n a n \_f l a g \leftarrow 1
\end{aligned}
\] \\
\hline \(\mathrm{V} \leftarrow+\) lnfinity & \(v \leftarrow S(p, s \mathrm{sc} 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow S(p, s \mathrm{sr} 2)\) & \(\mathrm{V} \leftarrow-\) Infinity & V \& Src2 & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src } 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow+\) lnfinity & \(v \leftarrow+\) lnfinity & \(v \leftarrow+\) lnfinity & \(v \leftarrow+\) lnfinity & \(\mathrm{V} \leftarrow+\) lnfinity & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{~d} \text { NaN } \\
& \text { vxisi_flag } \leftarrow 1
\end{aligned}
\] & v \& SrC2 & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow Q(\text { srč2) } \\
& \mathrm{vxsnnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \begin{tabular}{l}
\(v \leftarrow p\) \\
vxsnan_flag \(\leftarrow 1\)
\end{tabular} \\
\hline \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & V \& SrC2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 } 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline Explanation: & \\
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
\hline src2 & For \(\boldsymbol{x s n m s u b a d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT]. For \(\boldsymbol{x s n m s u b m d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XB]. \\
\hline src3 & For \(\boldsymbol{x s n m s u b a d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XB]. For \(\boldsymbol{x s n m s u b m d p}\), the double-precision floating-point value in doubleword element 0 of VSR[XT]. \\
\hline dQNaN & Default quiet NaN (0x7FF8_0000_0000_0000). \\
\hline NZF & Nonzero finite number. \\
\hline Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. \\
\hline \(Q(x)\) & Return a QNaN with the payload of x . \\
\hline \(\mathrm{S}(\mathrm{x}, \mathrm{y})\) & Return the normalized sum of floating-point value \(x\) and negated floating-point value \(y\), having unbounded range and precision. Note: If \(x=y, v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline \(\mathrm{M}(\mathrm{x}, \mathrm{y})\) & Return the normalized product of floating-point value \(x\) and floating-point value \(y\), having unbounded range and precision. \\
\hline p & The intermediate product having unbounded range and precision. \\
\hline v & The intermediate result having unbounded range and precision. \\
\hline
\end{tabular}

Table 96.Actions for xsnmsub(alm)dp

\section*{VSX Scalar Negative Multiply-Subtract Single-Precision XX3-form}
xsnmsubasp \(\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)

xsnmsubmsp \(\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & 11 & A & 16 & B & 21 & 153 & \[
\begin{aligned}
& |2 \times 8 \times X|>\mid \\
& 29303
\end{aligned}
\] \\
\hline
\end{tabular}
```

reset_xflags()
if "xsnmsubasp" then do
src1 \leftarrow VSR[32xAX+A].dword[0]
src2 }\leftarrow\operatorname{VSR[32xTX+T].dword[0]
src3 \leftarrow VSR[32xBX+B].dword[0]
end
if "xsnmsubmsp" then do
src1 \leftarrowVSR[32xAX+A].dword[0]
src2 }\leftarrow\textrm{VSR}[32\timesBX+B].dword[0
src3}\leftarrow\operatorname{VSR[32xTX+T].dword[0]
end
v }\leftarrow\mathrm{ MultiplyAddDP(src1,src3,NegateDP(src2)))
result }\leftarrow\operatorname{NegateSP(RoundToSP(RN,v))
if(vxsnan_flag) then SetFX(VXSNAN)
if(vximz_flag) then SetFX(VXIMZ)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(0X)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
vex_flag \leftarrow VE \& (vxsnan_flag | vximz_flag | vxisi_flag)
if( ~vex_flag ) then do
VSR[32xTX+T].dword[0] \leftarrow ConvertSPtoSP64 (result)
VSR[32xTX+T].dword[1] \& OxUUUU_UUUU_UUUU_UUUU
FPRF }\leftarrow\mathrm{ ClasSSP(result)
FR}\leftarrow\mathrm{ inc_flag
FI }\leftarrowxx_fla
end
else do
FR}\leftarrow0\textrm{ObO
FI }\leftarrow0\textrm{ObO
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

For xsnmsubasp, do the following.
- Let src1 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X A]\).
- Let src2 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X T]\).
- Let src3 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X B]\).

\section*{For xsnmsubmsp, do the following.}
- Let src1 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X A]\).
- Let src2 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X B]\).
- Let src3 be the double-precision floating-point value in doubleword element 0 of \(\operatorname{VSR}[X T]\).
src1 is multiplied \({ }^{[1]}\) by src3, producing a product having unbounded range and precision.

See part 1 of Table 97, "Actions for xsnmsub(a|m)sp," on page 626.
\(\operatorname{src} 2\) is negated and added \({ }^{[2]}\) to the product, producing a sum having unbounded range and precision.

The sum is normalized \({ }^{[3]}\).
See part 2 of Table 97, "Actions for xsnmsub(a|m)sp," on page 626.

The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is negated and placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, \(\mathrm{VSR}[\mathrm{XT}]\) and FPRF are not modified, and FR and FI are set to 0 .

See Table 93, "Scalar Floating-Point Final Result with Negation," on page 613.

\footnotetext{
1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( \(\mathrm{G}, \mathrm{R}\), and X ) enter into the computation.
3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}

\section*{Special Registers Altered}
FPRF FR FI FX OX UX XX VXSNAN VXISI VXIMZ

VSR Data Layout for xsnmsub(a|m)sp
src1 = VSR[XA]
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline src2 = \(\boldsymbol{x s n m s u b a s p}\) ? VSR[XT] : VSR[XB] \\
\hline DP & unused \\
\hline src3 = xsnmsubasp ? VSR[XB] : VSR[XT] \\
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\) VSR[XT]
\begin{tabular}{|l|l|}
\hline DP & undefined \\
\hline 0 & 64
\end{tabular}

\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src3} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow-\) Infinity & \(p \leftarrow-\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(p \leftarrow+\) Infinity & \(p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)\) & \(p \leftarrow \operatorname{srcl}\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \hline \mathrm{p} \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow+\) Zero & \(p \leftarrow+\) Zero & \(\mathrm{p} \leftarrow-\) Zero & \(\mathrm{p} \leftarrow-\) Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow\) src3 & \[
\begin{aligned}
& \mathrm{p} \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{p} \leftarrow-\) Zero & \(\mathrm{p} \leftarrow-\) Zero & \(\mathrm{p} \leftarrow+\) Zero & \(\mathrm{p} \leftarrow+\) Zero & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow\) Src3 & \[
\begin{aligned}
& p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)\) & \(p \leftarrow+\) Infinity & \(p \leftarrow \operatorname{src} 3\) & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src3 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(p \leftarrow-\) Infinity & \(p \leftarrow+\) Infinity & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow d Q N a N \\
& \text { vximz_flag } \leftarrow 1
\end{aligned}
\] & \(p \leftarrow+\) Infinity & \(p \leftarrow+\) Infinity & \(p \leftarrow\) src3 & \[
\begin{aligned}
& p \leftarrow Q(\text { src } 3) \\
& \text { vxsnan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \(p \leftarrow \operatorname{src} 1\) & \[
\begin{aligned}
& p \leftarrow \operatorname{src1} \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{p} \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& v x \text { snan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline p \leftarrow Q(\text { src1) } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{p} \leftarrow Q(\text { src1 } 1) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& p \leftarrow Q(\text { src1 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{c} 
Part 2: \\
Subtract \\
\hline \begin{tabular}{c}
-Infinity \\
\hline -NZF \\
\hline -Zero \\
\hline +Zero \\
\hline +NZF \\
\hline +Infinity \\
\hline \begin{tabular}{c} 
QNaN \& \\
src1 is a NaN
\end{tabular} \\
\hline \begin{tabular}{c} 
QNaN \& \\
src1 not a NaN
\end{tabular} \\
\hline
\end{tabular} \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src2} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow \mathrm{~d} Q \mathrm{NaN} \\
& \text { vxisi_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(\mathrm{V} \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow-\) Infinity & \(V \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline v \leftarrow Q(\text { src2 }) \\
& v x \text { snan_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] \\
\hline \(V \leftarrow+\) Infinity & \(v \leftarrow S(p, s r c 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow S\left(p, s r^{\prime} 2\right)\) & \(\checkmark \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{src} 2\) & \[
\begin{array}{|l}
\hline \mathrm{V} \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline \(V \leftarrow+\) Infinity & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(v \leftarrow\) Rezd & \(v \leftarrow\)-Zero & \(\mathrm{V} \leftarrow-\mathrm{src} 2\) & \(\checkmark \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow Q(\text { src2 }) \\
\text { vxsnan_flag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline \(V \leftarrow+\) Infinity & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(v \leftarrow+\) Zero & \(v \leftarrow\) Rezd & \(\mathrm{V} \leftarrow-\mathrm{src} 2\) & \(V \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow Q(\text { (src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(V \leftarrow+\) Infinity & \(v \leftarrow S(p, s r c 2)\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(\mathrm{v} \leftarrow \mathrm{S}(\mathrm{p}, \mathrm{sc} 22)\) & \(\mathrm{V} \leftarrow-\) Infinity & \(\mathrm{V} \leftarrow \mathrm{src} 2\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(v \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \(V \leftarrow+\) Infinity & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{~d} Q \mathrm{NaN} \\
& \text { vxisi_flag } \leftarrow 1 \\
& \hline
\end{aligned}
\] & \(\mathrm{V} \leftarrow \mathrm{src} 2\) & \[
\begin{aligned}
& v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \[
\begin{aligned}
& v \leftarrow p \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(v \leftarrow p\) & \(\mathrm{V} \leftarrow \mathrm{SrC2}\) & \[
\begin{aligned}
& v \leftarrow Q(\text { src2 }) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}

\section*{Explanation:}
\begin{tabular}{|c|c|}
\hline src1 & The double-precision floating-point value in VSR[XA].dword[0]. \\
\hline src2 & For xsnmsubasp, the double-precision floating-point value in VSR[XT].dword[0]. For \(\boldsymbol{x s n m s u b m s p}\), the double-precision floating-point value in VSR[XB].dword[0]. \\
\hline src3 & For \(\boldsymbol{x s n m s u b a s p}\), the double-precision floating-point value in VSR[XB].dword[0]. For \(\boldsymbol{x s n m s u b m s p}\), the double-precision floating-point value in VSR[XT].dword[0]. \\
\hline dQNaN & Default quiet NaN (0x7FF8_0000_0000_0000). \\
\hline NZF & Nonzero finite number. \\
\hline Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. \\
\hline \(Q(x)\) & Return a QNaN with the payload of x . \\
\hline \(S(x, y)\) & Return the normalized sum of floating-point value \(x\) and negated floating-point value \(y\), having unbounded range and precision. Note: If \(x=y, v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline M ( \(\mathrm{x}, \mathrm{y}\) ) & Return the normalized product of floating-point value \(x\) and floating-point value y , having unbounded range and precision. \\
\hline p & The intermediate product having unbounded range and precision. \\
\hline v & The intermediate result having unbounded range and precision. \\
\hline
\end{tabular}

Table 97.Actions for xsnmsub(alm)sp

\section*{VSX Scalar Negative Multiply-Subtract Quad-Precision [using round to Odd] X-form}
```

lll

| 063 | ${ }_{6} \text { VRT }$ | ${ }_{11} \text { VRA }$ | ${ }_{16} \text { VRB }$ | 21 | 484 | RO |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |

```
```

if MSR.VSX=0 then VSX_Unavailable(l

```
if MSR.VSX=0 then VSX_Unavailable(l
reset_xflags()
reset_xflags()
srcl &bfp_CONVERT_FROM_BFP128(VSR[VRA+32])
srcl &bfp_CONVERT_FROM_BFP128(VSR[VRA+32])
src2 \leftarrowbfp_CONVERT_FROM_BFP128(VSR[VRT+32])
src2 \leftarrowbfp_CONVERT_FROM_BFP128(VSR[VRT+32])
srC3 &bfp_CONVERT_FROM_BFP128(VSR[VRB+32])
srC3 &bfp_CONVERT_FROM_BFP128(VSR[VRB+32])
v &bfp_MULTIPLY_ADD(srcl, src3, bfp_NEGATE(src2))
v &bfp_MULTIPLY_ADD(srcl, src3, bfp_NEGATE(src2))
rnd < bfp_NEGATE(bfp_ROUND_TO_BFP128(RO, FPSCR, RN, v))
rnd < bfp_NEGATE(bfp_ROUND_TO_BFP128(RO, FPSCR, RN, v))
result & bfp_CONVERT_TO_BFP128(rnd)
result & bfp_CONVERT_TO_BFP128(rnd)
if(vxsnan_flag) then SetFX(FPSCR.VXSNAN)
if(vxsnan_flag) then SetFX(FPSCR.VXSNAN)
if(vximz_flag) then SetFX(FPSCR.VXIMZ)
if(vximz_flag) then SetFX(FPSCR.VXIMZ)
if(Ox_flag) then SetFX(FPSCR.OX)
if(Ox_flag) then SetFX(FPSCR.OX)
if(ux_flag) then SetFX(FPSCR.UX)
if(ux_flag) then SetFX(FPSCR.UX)
if(xx_flag) then SetFX(FPSCR.XX)
if(xx_flag) then SetFX(FPSCR.XX)
vx_flag \leftarrowvxsnan_flag| vxim_flag
vx_flag \leftarrowvxsnan_flag| vxim_flag
ex_flag \leftarrowFPSCR.VE & VX_flag
ex_flag \leftarrowFPSCR.VE & VX_flag
if ex_flag=0 then do
if ex_flag=0 then do
    VSR[[VRT+32]}\leftarrow\mathrm{ result
    VSR[[VRT+32]}\leftarrow\mathrm{ result
    FPSCR,FPRF \leftarrowfprf_CLASS_BFP128(result)
    FPSCR,FPRF \leftarrowfprf_CLASS_BFP128(result)
end
end
FPSCR,FR\leftarrow(vx_flag=0) & inc_flag
FPSCR,FR\leftarrow(vx_flag=0) & inc_flag
FPSCR,Fl \leftarrow(vx_flag=0) & xx_flag
```

FPSCR,Fl \leftarrow(vx_flag=0) \& xx_flag

```
(RO=1)

Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let src2 be the floating-point value in VSR[VRT+32] represented in quad-precision format.

Let src3 be the floating-point value in VSR[VRB+32] represented in quad-precision format.

If either \(\operatorname{srcl}, \operatorname{src} 2\), or \(\operatorname{srcs}\) is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN is set to 1.

If \(\operatorname{srcl}\) is an Infinity value and \(\operatorname{srcs}^{\text {is a Zero value, or if }}\) srcl is a Zero value and srcs is an Infinity value, an Invalid Operation exception occurs and VXI MZ is set to 1.

If \(\mathrm{srCl}_{2}\) and the product of \(\mathrm{srCl}^{2}\) and \(\mathrm{srCl}^{\text {are Infinity }}\) values having same signs, an Invalid Operation exception occurs and VXISI is set to 1 .

If srcl is a Signalling NaN , the result is the Quiet NaN corresponding to sicl.

Otherwise, if srcl is a Quiet NaN , the result is srcl .
Otherwise, if \(\mathrm{SrC2}\) is a Signalling NaN , the result is the Quiet NaN corresponding to sic2.

Otherwise, if \(\mathrm{src2}\) is a Quiet NaN , the result is \(\mathrm{src2}\).
Otherwise, if \(\mathrm{SrCl}^{\mathrm{Cl}}\) is a Signalling NaN , the result is the Quiet NaN corresponding to src .

Otherwise, if \(\mathrm{src}_{\mathrm{Cl}}\) is a Quiet NaN , the result is \(\mathrm{src3}\).
Otherwise, if srCl is an Infinity value and srcs is a Zero value, or if \(\operatorname{srcl}\) is a Zero value and \(\operatorname{srcs}\) is an Infinity value, the result is the default Quiet \(\mathrm{NaN}^{[1]}\).

Otherwise, if the product of \(\operatorname{srcl}\) and \(\operatorname{srcs}\), and \(\operatorname{srcz}\) are Infinity values having same signs, the result is the default Quiet NaN .

Otherwise, do the following.
srcl is multiplied by \(\operatorname{src} 3\), producing a product having unbounded significand precision and exponent range.

See part 1 of Table 88. "Actions for xsmsubqp[o]".
srct is negated and added to the product, producing a sum having unbounded range and precision.

See part 2 of Table 88. "Actions for xsmsubqp[o]".
If the intermediate result is Tiny (i.e., the unbiased exponent is less than -16382) and \(U E=0\), the significand is shifted right \(N\) bits, where \(N\) is the difference between \(\cdot 16382\) and the unbiased exponent of the intermediate result. The exponent of the intermediate result is set to the value - 16382.

If \(R O=1\), let the rounding mode be Round to Odd. Otherwise, let the rounding mode be specified by RN. Unless the result is an Infinity or a Zero, the intermediate result is rounded to quad-precision using the specified rounding mode.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is negated and placed into VSR[VRT+32] in quad-precision format.

FPRF is set to the class and sign of the result. FR is set to indicate if the rounded result was incremented. FI is set to indicate the result is inexact.

If a trap-disabled Invalid Operation exception occurs, \(F R\) and \(F I\) are set to 0 .

\section*{Version 3.0}

If a trap-enabled Invalid Operation exception occurs, VSR[VRT+32] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

Special Registers Altered:
FPRF FR FI
FX VXSNAN VXIMZ VXISI OX UX XX
| VSR Data Layout for xsnmsubqp[o]
I VSR[VRA+32]
I \(\operatorname{src1}\)
\| VSR[VRT+32]
I \(\operatorname{src2}\)
I VSR[VRB+32]
1 sic3
I VSR[VRT+32]
I \(\quad\) tgt


\section*{\| Table 98.Actions for xsnmsubqp[o]}

\section*{Version 3.0}

VSX Scalar Round to Double-Precision Integer using round to Nearest Away XX2-form


Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src is rounded to an integer using the rounding mode Round to Nearest Away.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result. FR is set to 0 . Fl is set to 0 .

If a trap-enabled invalid operation exception occurs, \(\mathrm{VSR}[\mathrm{XT}]\) and FPRF are not modified, and FR and FI are set to 0 .

\section*{Special Registers Altered}
FPRF FR=0b0 FI=0b0 FX VXSNAN

VSR Data Layout for xsrdpi
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)


\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

\section*{VSX Scalar Round to Double-Precision Integer exact using Current rounding mode XX2-form}
```

xsrdpic XT,XB

| 60 | T | III |  | B |  | 107 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 6 | 11 | 16 |  | 21 |  |


| $X T$ | $\leftarrow T X \\| T$ |
| :--- | :--- | :--- |
| $X B$ | $\leftarrow B X \\| B$ |

reset xflag
reset_xflags()
src }\leftarrow\textrm{VSR}[XB]{0:63
if(RN=0b00) then result{0:63} \leftarrow RoundToDPIntegerNearEven(src)
if(RN=0b01) then result{0:63}}\leftarrow\mathrm{ RoundToDPIntegerTrunc(src)
if(RN=0b10) then result{0:63} \leftarrowRoundToDPIntegerCeil(src)
if(RN=0b11) then result{0:63}}\leftarrow\mathrm{ RoundToDPIntegerFloor(src)
if(vxsnan_flag) then SetFX(VXSNAN)
if(xx_flag) then SetFX(XX)
vex_flag \leftarrow VE \& vxsnan_flag
if( ~vex_flag ) then do
VSR[XT] \leftarrow result || 0xUUUU_UUUU_UUUU_UUUU
FPRF}\leftarrow\mathrm{ ClassDP(result)
FR}\leftarrow\mathrm{ inc_flag
FI }\leftarrowxx_fla
end
else do
FR}\leftarrow00b
FI }\leftarrow0\textrm{b}
end

```

Let XT be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).

Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src is rounded to an integer using the rounding mode specified by RN.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. FI is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

\section*{Special Registers Altered}
```

            FPRF FR FI FX XXVXSNAN
    ```

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

VSR Data Layout for xsrdpic
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|r|l|}
\hline DP & unused \\
\hline tgt \(=\) VSR[XT] \\
\hline DP & undefined \\
\hline 0 & 64 \\
\hline
\end{tabular}

VSX Scalar Round to Double-Precision Integer using round toward -Infinity XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|l|}{xsrdpim XT, XB} \\
\hline \[
\begin{array}{ll} 
& 60 \\
0
\end{array}
\] & \[
\begin{array}{ll} 
& \mathrm{T} \\
6 &
\end{array}
\] & \[
\int_{11} \quad \text { III }
\] & \[
1_{16} B
\] & 21 & 121 & \(|\mathrm{BX} \times \mathrm{TX}|\) \\
\hline \multicolumn{7}{|l|}{XT \(\quad \leftarrow \mathrm{TX} \|\) T} \\
\hline \multicolumn{7}{|l|}{} \\
\hline \multicolumn{7}{|l|}{reset_xflags()} \\
\hline \multicolumn{7}{|l|}{result \(\{0: 63\} \leftarrow\) RoundToDPIntegerFloor (VSR[XB] \(00: 63\}\) )} \\
\hline \multicolumn{7}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{7}{|l|}{FR \(\quad \leftarrow 0\) b0} \\
\hline \multicolumn{7}{|l|}{FI} \\
\hline \multicolumn{7}{|l|}{vex_flag \(\leftarrow\) VE \& vxsnan_flag} \\
\hline \multicolumn{7}{|l|}{if( ~vex_flag ) then do} \\
\hline \multicolumn{7}{|c|}{VSR[XT] \(\leftarrow\) result || 0xUUUU_UUUU_UUUU_UUUU} \\
\hline \multicolumn{7}{|c|}{FPRF \(\leftarrow\) ClassDP(result)} \\
\hline \multicolumn{7}{|l|}{end} \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src is rounded to an integer using the rounding mode Round toward -Infinity.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to 0 . Fl is set to 0 .

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .
```

Special Registers Altered
FPRF FR=0b0 FI=0b0 FX VXSNAN

```

VSR Data Layout for xsrdpim
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline DP & \multicolumn{2}{|c|}{ undefined } \\
\hline 0 & 64
\end{tabular}

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

VSX Scalar Round to Double-Precision Integer using round toward +Infinity XX2-form
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{xsrdpip XT, XB} \\
\hline \[
{ }_{0} 60
\] & \({ }_{6}{ }^{\text {T }}\) & \[
{ }_{11} \quad \text { III }
\] & \[
{ }_{16} \quad \text { B }
\] & \[
\begin{array}{ll} 
& 105 \\
21 &
\end{array}
\] & BX \(\begin{gathered}\text { TX } \\ 30 \\ 31\end{gathered}\) \\
\hline
\end{tabular}
\begin{tabular}{lr} 
XT & \(\leftarrow T X \| T\) \\
XB & \(\leftarrow B X \| B\)
\end{tabular}
\(X B \quad \leftarrow B X \| B\)
reset_xflags()
result \(\{0: 63\} \leftarrow\) RoundToDPIntegerCeil(VSR[XB]\{0:63\})
if(vxsnan_flag) then SetFX(VXSNAN)
FR \(\leftarrow 0\) bo
\(\mathrm{FI} \quad \leftarrow 0 \mathrm{~b} 0\)
vex_flag \(\leftarrow\) VE \& vxsnan_flag
if( ~vex_flag ) then do
VSR[XT] \(\leftarrow\) result || 0xUUUU_UUUU_UUUU_UUUU
FPRF \(\leftarrow\) ClassDP(result)
end
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src is rounded to an integer using the rounding mode Round toward + Infinity.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result. FR is set to 0 . Fl is set to 0 .

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

\section*{Special Registers Altered}

FPRF FR=0b0 FI=0b0 FX VXSNAN
VSR Data Layout for xsrdpip
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline DP & \multicolumn{2}{c|}{ undefined } \\
\hline 0 & 64
\end{tabular}

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

VSX Scalar Round to Double-Precision Integer using round toward Zero XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|l|}{xsrdpiz XT,XB} \\
\hline \[
\begin{array}{ll} 
& 60 \\
0
\end{array}
\] & \[
{ }_{6} \quad \mathrm{~T}
\] & \({ }_{11}\) III & 16 & & 89 & BXX \(\mid\) TX \\
\hline \multicolumn{7}{|l|}{XT \(\quad \leftarrow T X \| T\)} \\
\hline \multicolumn{7}{|l|}{\(X B \quad \leftarrow B X \| B\)} \\
\hline \multicolumn{7}{|l|}{reset_xflags()} \\
\hline \multicolumn{7}{|l|}{result\{0:63\} \(\leftarrow\) RoundToDPIntegerTrunc(VSR[XB]\{0:63\})} \\
\hline \multicolumn{7}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{7}{|l|}{FR \(\quad \leftarrow 0\) b0} \\
\hline \multicolumn{7}{|l|}{FI} \\
\hline \multicolumn{7}{|l|}{vex_flag \(\leftarrow\) VE \& vxsnan_flag} \\
\hline \multicolumn{7}{|l|}{if( ~vex_flag ) then do} \\
\hline \multicolumn{7}{|l|}{VSR[XT] \(\leftarrow\) result || 0xUUUU_UUUU_UUUU_UUUU} \\
\hline \multicolumn{7}{|l|}{FPRF \(\quad \leftarrow\) ClassDP(result)} \\
\hline end & & & & & & \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src is rounded to an integer using the rounding mode Round toward Zero.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to 0 . Fl is set to 0 .

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

\section*{Special Registers Altered}

FPRF FR=0b0 FI=0b0 FX VXSNAN

VSR Data Layout for xsrdpiz
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)


\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

VSX Scalar Reciprocal Estimate Double-Precision XX2-form
```

xsredp XT,XB

| 60 |  | T | III |  | B |  | 90 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  | 16 |  |
| 21 |  |  | $3 \times T X$ |  |  |  |  |
| 3031 |  |  |  |  |  |  |  |

$X T \quad \leftarrow T X \| T$
XB }\leftarrow\textrm{BX|B
reset_xflags()
v{0:inf}}\leftarrow\mathrm{ ReciprocalEstimateDP(VSR[XB]{0:63})
result{0:63} \leftarrow RoundToDP(RN,v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(zx_flag) then SetFX(ZX)
vex_flag }\leftarrow\mathrm{ VE \& vxsnan_flag
zex_flag}\leftarrow\mathrm{ ZE \& zx_flag
if( ~vex_flag \& ~zex_flag ) then do
VSR[XT] \leftarrow result || 0xUUUU_UUUU_UUUU_UUUU
FPRF}\leftarrow\mathrm{ ClassDP(result)
FR}\leftarrow0\mathrm{ bU
FI}\leftarrow0\textrm{Ob
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].

A double-precision floating-point estimate of the reciprocal of src is placed into doubleword element 0 of VSR[XT] in double-precision format.

Unless the reciprocal of src would be a zero, an infinity, or a QNaN, the estimate has a relative error in precision no greater than one part in 16384 of the reciprocal of src. That is,
\[
\left|\frac{\text { estimate }-\frac{1}{s r c}}{\frac{1}{s r c}}\right| \leq \frac{1}{16384}
\]

Operation with various special values of the operand is summarized below.
\begin{tabular}{ccc}
\hline Source Value & Result & Exception \\
\hline -Infinity & -Zero & None \\
\hline -Zero & -Infinity \(^{1}\) & ZX \\
\hline +Zero & +lnfinity \(^{1}\) & ZX \\
\hline +Infinity & +Zero & None \\
\hline SNaN & QNaN \(^{2}\) & VXSNAN \\
\hline QNaN & QNaN & None \\
\hline
\end{tabular}
1. No result if \(\mathrm{ZE}=1\).
2. No result if \(\mathrm{VE}=1\).

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result. FR is set to an undefined value. FI is set to an undefined value.

If a trap-enabled invalid operation exception or a trap-enabled zero divide exception occurs, VSR[XT] and FPRF are not modified.

The results of executing this instruction is permitted to vary between implementations, and between different executions on the same implementation.

\section*{Special Registers Altered}

FPRF FR=0bU FI=0bU FX OX UX
\(X X=0 b U\) VXSNAN
VSR Data Layout for xsredp
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline DP & \multicolumn{1}{c|}{ undefined } \\
\hline 0 & 64
\end{tabular}

\section*{VSX Scalar Reciprocal Estimate Single-Precision XX2-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{xsresp XT,XB} \\
\hline \[
\begin{array}{ll} 
& 60 \\
0 &
\end{array}
\] & \({ }_{6}{ }^{\text {T }}\) & \[
\int_{11} I I I
\] & \[
{ }_{16} \quad \text { B }
\] & \[
26
\] & \(|\)\begin{tabular}{l|l|} 
BXX \\
30 & 31 \\
\end{tabular} \\
\hline
\end{tabular}
```

reset_xflags()
src }\leftarrow\operatorname{VSR[32\timesBX+B].dword[0]
v}\leftarrow\mathrm{ ReciprocalEstimateDP(src)
result }\leftarrow\mathrm{ RoundTOSP(RN,v)

```
if (vxsnan_flag) then SetFX (VXSNAN)
if (ox_flag) then SetFX(OX)
if (ux_flag) then SetFX(UX)
if (ObU) then SetFX(XX)
if(zx_flag) then \(\operatorname{SetFX}(\mathrm{ZX})\)
vex_flag \(\leftarrow\) VE \& vxsnan_flag
zex_flag \(\leftarrow\) ZE \& zx_flag
if ( ~vex_flag \& ~zex_flag ) then do
    VSR[32xTX+T].dword[0] \(\leftarrow\) ConvertSPtoSP64 (result)
    VSR [32xTX+T] .dword[1] \(\leftarrow\) 0xUUUU_UUUU_UUUU_UUUU
    FPRF \(\leftarrow\) ClassSP (result)
    \(\mathrm{FR} \leftarrow \mathrm{ObU}\)
    \(\mathrm{FI} \leftarrow \mathrm{ObU}\)
end
else do
    \(\mathrm{FR} \leftarrow \mathrm{ObO}\)
    \(\mathrm{FI} \leftarrow 0 \mathrm{bO}\)
end

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

A single-precision floating-point estimate of the reciprocal of src is placed into doubleword element 0 of VSR[ XT] in double-precision format.

Unless the reciprocal of sic would be a zero, an infinity, the result of a trap-disabled Overflow exception, or a QNaN, the estimate has a relative error in precision no greater than one part in 16384 of the reciprocal of src . That is,
\[
\left|\frac{\text { estimate }-\frac{1}{\operatorname{src}}}{\frac{1}{\mathrm{src}}}\right| \leq \frac{1}{16384}
\]

Operation with various special values of the operand is summarized below.
\begin{tabular}{ccc}
\hline Source Value & Result & Exception \\
\hline -Infinity & -Zero & None \\
\hline -Zero & - Infinity \(^{1}\) & ZX \\
\hline +Zero & + \(^{\text {Infinity }}{ }^{1}\) & ZX \\
\hline +Infinity & +Zero & None \\
\hline SNaN & QNaN \(^{2}\) & VXSNAN \\
\hline QNaN & QNaN & None \\
\hline
\end{tabular}
1. No result if \(Z E=1\).
2. No result if \(\mathrm{VE}=1\).

The contents of doubleword element 1 of VSR[ XT] are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to an undefined value. Fl is set to an undefined value.

If a trap-enabled invalid operation exception or a trap-enabled zero divide exception occurs, VSR[ XT] and FPRF are not modified.

The results of executing this instruction is permitted to vary between implementations, and between different executions on the same implementation.

\section*{Special Registers Altered}
```

FPRF FR=ObU FI=ObU FX OX UX ZX XX=ObU
VXSNAN

```

VSR Data Layout for xsresp
src = VSR[XB]
\begin{tabular}{|c|c|}
\hline \multicolumn{1}{|c|}{ DP } & unused \\
\hline tgt \(=\) VSR[XT] & \\
\hline DP & undefined \\
\hline 0 & 64
\end{tabular}

VSX Scalar Round to Quad-Precision Integer [with Inexact] Z23-form
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline xsrqpi & \multicolumn{6}{|c|}{R,VRT,VRB, RMC} & \multicolumn{2}{|r|}{\multirow[t]{2}{*}{\[
\begin{aligned}
& (E X=0) \\
& (E X=1)
\end{aligned}
\]}} \\
\hline xsrqpix & & , VRT & R & B,RMC & & & & \\
\hline \[
\begin{array}{ll} 
& 63 \\
0 & \\
\hline
\end{array}
\] & \[
6
\] & 111 & \begin{tabular}{|c|}
\(R\) \\
15 \\
16
\end{tabular} & VRB & RMC
21 & 23 & 5 & EX \\
\hline
\end{tabular}
if MSR. VSX=O then VSX_Unavailable ()
reset_xflags()
if \(R=0\) then do
if RMC=0b00 then II Round to Nearest Away rmode \(\leftarrow 0 b 100\)
if RMC=0b11 then do
if FPSCR. RN=ObOO then || Round to Nearest Even rmode \(\leftarrow 0 b 000\)
if FPSCR, RN=Ob01 then I/ Round towards Zero
rmode \(\leftarrow 06001\)
if FPSCR. RN=Ob10 then /| Round towards + Infinity rmode \(\leftarrow 06010\)
if FPSCR. RN=Ob11 then I/ Round towards - Infinity rmode \(\leftarrow 06011\)
end
end
else do || \(R=1\)
if RMC=0000 then
rmode \(\leftarrow 0 b 000\)
if RMC=0bo1 then
rmode \(\leftarrow 0 b 001\)
if RMC=Ob10 then II Round towards thninity
\(r\) mode \(\leftarrow\) ob010
if RMC=0b11 then II Round towards - Infinity rmode \(\leftarrow 0 b 011\)
end
sic \(\leftarrow\) bf P_CONVERT_RROM_BFP128(VSR[ VRB +32\(]\) )
if src.class. SNaN then do
result \(\leftarrow\) bfp_CONVERT_TO_bFP128(bfp_QUIET(stc))
vxsnan_flag \(\leftarrow 1\)
end
else if src.class. QNaN
src.class.Infinity
scc.class.Zero then
result \(\leftarrow\) bfp_CONERT_TO_bFP128(src)
else do
and \(\&\) bfp_round_tol INTEGER( rmode, sic)
result \(\leftarrow\) bf _ CONVERT_TO_bFP128(rnd)
end
if(uxsnan_flag) then SetFX(fPSCR. vxswan)
if(xx_flag \& EX) then SetFX(FPSCR. XX)
ex_flag \(\leftarrow\) FPSCR.VE \(\&\) vxsnan_flag
if ex filag=0 then do
VSR[ [VRT +32 ] \(\leftarrow\) result
FPSCR.FPRF \(\leftarrow\) fpri_Class_bPP128(result)
end
FPSCR. \(F R \leftarrow E X \&(\) uxsnan_ \(f l a g=0) \&\) inc_flag


Let \(R\) and RMC specify the rounding mode as follows.
\begin{tabular}{|c|c|c|c|}
\hline ¢ & \[
\sum_{\mathbb{X}}^{U}
\] &  & Rounding Mode \\
\hline 0 & 00 & - & Round to Nearest Away \\
\hline 0 & 01 & - & reserved \\
\hline 0 & 10 & - & reserved \\
\hline 0 & 11 & 00 & Round to Nearest Even \\
\hline 0 & 11 & 01 & Round towards Zero \\
\hline 0 & 11 & 10 & Round towards +Infinity \\
\hline 0 & 11 & 11 & Round towards -Infinity \\
\hline 1 & 00 & - & Round to Nearest Even \\
\hline 1 & 01 & - & Round towards Zero \\
\hline 1 & 10 & - & Round towards +Infinity \\
\hline 1 & 11 & - & Round towards -Infinity \\
\hline
\end{tabular}

Let \(\operatorname{sic}\) be the floating-point value in VSR[VRB +32 ] represented in quad-precision format.

If \(\operatorname{src}\) is a Signalling NaN , an Invalid Operation exception occurs, VXSNAN is set to 1 , and the result is the Quiet NaN corresponding to the Signalling NaN .

Otherwise, if \(s r c\) is a Quiet NaN, an Infinity, or a Zero, then the result is src .

Otherwise, src is rounded to an integer using the rounding mode rmode.

The result is placed into VSR[ VRT +32] in quad-precision format.

FPRF is set to the class and sign of the result.
For xsrqpi, \(F R\) is set to \(0, F I\) is set to 0 , and \(X X\) is not set by an Inexact exception.

For xsrqpix, \(F R\) is set to indicate if the result was incremented when rounded, Fl is set to indicate the result is inexact, and \(X X\) is set by an Inexact exception.

If a trap-disabled Invalid Operation exception occurs, FPRF is set to an undefined value.

If a trap-enabled Invalid Operation exception occurs, VSR[ VRT +32 ] and FPRF are not modified.

\section*{Special Registers Altered:}

FPRF VXSNAN FX



VSX Scalar Round Quad-Precision to Double-Extended Precision Z23-form
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{xsrqpxp R,VRT,VRB,RMC} \\
\hline \[
63
\] & \[
6_{6} \text { VRT }
\] & \(111 /\) & \(\left\lvert\,\)\begin{tabular}{l|l}
\(R\) \\
15
\end{tabular} 16 VRB\right. & \begin{tabular}{|l|} 
RMC \\
21
\end{tabular} & 23 & 37 & \begin{tabular}{|r}
1 \\
31
\end{tabular} \\
\hline
\end{tabular}
if MSR. VSX=0 then VSX_Unavailable()
reset_xflags()
if \(R=0\) then do
if RMC=ObOO the
II Round to Nearest Away
rmode \(\leftarrow 0 b 100\)
if RMC=Ob11 then do
if \(\operatorname{FPSCR}\). RN=ObOO then || Round to Nearest Even
rmode \(\leftarrow\) Obooo
if FPSCR. RN=ObO1 then I/ Round towards Zero
rmode \(\leftarrow 0 b 001\)
if FPSCR. RN=Ob10 then II Round towards +Infinity
rmode \(\leftarrow 06010\)
if FPSCR. RN=Ob11 then I/ Round towards.Infinity rmode \(\leftarrow 0 b 011\)
end
end
else do || \(R=1\)
if RMC=0600 the
rmode \(\leftarrow 06000\)
if RMC=0bo1 then
rmode \(\leftarrow 06001\)
if RMC=0b10 then \(r\) mode \(\leftarrow 0 b 010\)
if RMC=0b11 then rmode \(\leftarrow 0 b 011\)
end
sic \(\leftarrow\) bfp_Convert_from_bfp128/ (VSR[VBb+32])
ind \(\leftarrow\) bf p_round tö_bfp8o(rmode, stc)

if(uxsnan_flag) then SetFX(FPSCR.VXSNAN)
if(ox_flag) then SetFX(FPSCR. OX)
if(ux_flag) then SetFX(FPSCR.UX)
if(xx_flag) then SetFX(FPSCR. XX)
ex_flag \(\leftarrow\) FPSCR. VE \(\&\) vxsnan_flag
if ex_flag=0 then do
VSR[VRT +32\(] \leftarrow\) result
FPSCR.FPRF \(\leftarrow\) fprf_CLASS_BPP128(result)
end
FPSCR. FR \(\leftarrow(\) vxsnan_ \(f l a g=0) \& i n c \_\)flag
FPSCR. \(F 1 \leftarrow(\) vxs nan_flag \(=0) \& x x_{1} f\) fag

Let \(R\) and RMC specify the rounding mode as follows.
\begin{tabular}{|c|c|c|c|}
\hline ¢ & \[
\sum_{\mathbb{X}}^{U}
\] &  & Rounding Mode \\
\hline 0 & 00 & - & Round to Nearest Away \\
\hline 0 & 01 & - & reserved \\
\hline 0 & 10 & - & reserved \\
\hline 0 & 11 & 00 & Round to Nearest Even \\
\hline 0 & 11 & 01 & Round to Zero \\
\hline 0 & 11 & 10 & Round to +Infinity \\
\hline 0 & 11 & 11 & Round to -Infinity \\
\hline 1 & 00 & - & Round to Nearest Even \\
\hline 1 & 01 & - & Round to Zero \\
\hline 1 & 10 & - & Round to +Infinity \\
\hline 1 & 11 & - & Round to -Infinity \\
\hline
\end{tabular}

Let \(\operatorname{sic}\) be the floating-point value in VSR[VRB +32 ] represented in quad-precision format.

If \(\operatorname{src}\) is a Signalling NaN , an Invalid Operation exception occurs, VXSNAN is set to 1 , and the result is the Quiet NaN corresponding to the Signalling NaN , with the significand truncated to double-extended-precision.

Otherwise, if src is a Quiet NaN , then the result is src with the significand truncated to double-extended-precision.

Otherwise, if src is an Infinity or a Zero, the result is src.

Otherwise, src is rounded to double-extended precision (i.e., 15-bit exponent range and 64-bit significand precision) using the specified rounding mode.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into VSR[ VRT +32] in quad-precision format.

FPRF is set to the class and sign of the result. \(F R\) is set to indicate if the rounded result was incremented. Fl is set to indicate the result is inexact.

If a trap-disabled Invalid Operation exception occurs, \(F P R F\) is set to an undefined value, and FR and FI are set to 0 .

If a trap-enabled Invalid Operation exception occurs, VSR[VRT+32] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.
\begin{tabular}{|c|}
\hline Special Registers Altered: \\
\hline VSR Data Layout for xsrqpxp \\
\hline VSR[ VRB+32] \\
\hline sic \\
\hline VSR[ VRT +32] \\
\hline tgt \\
\hline
\end{tabular}

\section*{Version 3.0}

VSX Scalar Round to Single-Precision XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{xsrsp \(\quad \mathrm{XT}, \mathrm{XB}\)} \\
\hline 060 & \({ }_{6}{ }^{\text {T }}\) & \[
{ }_{11} \quad \text { III }
\] & 16 & B & 21 & 281 & BXXTX 3031 \\
\hline
\end{tabular}

\section*{reset_xflags()}
\(\operatorname{src} \quad \leftarrow \operatorname{VSR}[32 \times B X+B]\).dword[0]
result \(\leftarrow\) RoundToSP (RN, src)
if (vxsnan_flag) then SetFX (VXSNAN)
if(ox_flag) then SetFX(0X)
if (ux_flag) then SetFX(UX)
if (xx_flag) then \(\operatorname{SetFX}(X X)\)
vex_flag \(\leftarrow\) VE \& vxsnan_flag
if( ~vex_flag) then do \(\operatorname{VSR}[32 \times T X+T]\).dword \([0] \leftarrow\) ConvertSPtoSP64 (result) VSR [32×TX+T].dword[1] \(\leftarrow\) 0xUUUU_UUUU_UUUU_UUUU FPRF \(\leftarrow\) ClassSP (result)
\(\mathrm{FR} \leftarrow\) inc_flag FI \(\leftarrow x x \_f l a g\)
end
else do
\(\mathrm{FR} \leftarrow \mathrm{ObO}\)
\(\mathrm{FI} \leftarrow \mathrm{ObO}\)
end

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of \(\operatorname{VSR}[\mathrm{XT}]\) are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified.

\section*{Special Registers Altered}

FPRF FR FI FX OX UX XX VXSNAN

\section*{VSX Scalar Reciprocal Square Root Estimate Double-Precision XX2-form}
xsrsqrtedp \(\quad \mathrm{XT}, \mathrm{XB}\)
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline -60 & 6 & T & 11 & III & & B & 1 & 74 & BXTX \\
\hline
\end{tabular}

if( ~vex_flag \& ~zex_flag ) then do
VSR[XT] \(\leftarrow\) result || 0xUUUU_UUUU_UUUU_UUUU
FPRF \(\leftarrow\) ClassDP(result)
\(\mathrm{FR} \leftarrow 0 \mathrm{bU}\)
\(\mathrm{FI} \leftarrow 0 \mathrm{bU}\)
end
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].

A double-precision floating-point estimate of the reciprocal square root of src is placed into doubleword element 0 of VSR[XT] in double-precision format.

Unless the reciprocal of the square root of src would be a zero, an infinity, or a QNaN, the estimate has a relative error in precision no greater than one part in 16384 of the reciprocal of the square root of src. That is,
\[
\left|\frac{\text { estimate }-\frac{1}{\sqrt{s r c}}}{\frac{1}{\sqrt{\mathrm{src}}}}\right| \leq \frac{1}{16384}
\]

Operation with various special values of the operand is summarized below.
\begin{tabular}{ccc}
\hline Source Value & Result & Exception \\
\hline -Infinity & QNaN \(^{1}\) & VXSQRT \\
\hline -Finite & QNaN \(^{1}\) & VXSQRT \\
\hline -Zero & -Infinity \(^{2}\) & ZX \\
\hline +Zero & + \(^{\text {Infinity }}{ }^{2}\) & ZX \\
\hline +Infinity & + Zero & None \\
\hline SNaN & QNaN \(^{1}\) & VXSNAN \\
\hline QNaN & QNaN & None \\
\hline
\end{tabular}
1. No result if \(\mathrm{VE}=1\).
2. No result if \(\mathrm{ZE}=1\).

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to an undefined value. FI is set to an undefined value.

If a trap-enabled invalid operation exception or a trap-enabled zero divide exception occurs, VSR[XT] and FPRF are not modified.

The results of executing this instruction is permitted to vary between implementations, and between different executions on the same implementation.

\section*{Special Registers Altered}
```

    FPRF FR=0bU FI=0bU FX
    ```
    XX=0bU VXSNAN VXSQRT

\section*{VSR Data Layout for xsrsqrtedp}
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\operatorname{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline DP & \multicolumn{1}{c|}{ undefined } \\
\hline 0 & 64 \\
\hline
\end{tabular}

VSX Scalar Reciprocal Square Root Estimate Single-Precision XX2-form
\begin{tabular}{|c|c|c|c|c|c|}
\hline xsrsqrt & & XT, XB & & & \\
\hline \[
\begin{array}{ll} 
& 60 \\
0
\end{array}
\] & \[
{ }_{6} \quad \mathrm{~T}
\] & 111 III & 16 & 210 & BX \(\begin{aligned} & \text { TXX } \\ & 30 \\ & 31\end{aligned}\) \\
\hline
\end{tabular}
```

reset_xflags()
SrC }\leftarrow\mathrm{ VSR[32xBX+B].dword[0]
v}\quad\leftarrow\mathrm{ ReciprocalSquareRootEstimateDP(src)
result }\leftarrow\mathrm{ RoundToSP(RN,v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxsqrt_flag) then SetFX(VXSQRT)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(0bU) then SetFX(XX)
if(zx_flag) then SetFX(ZX)
vex_flag \leftarrow VE \& (vxsnan_flag | vxsqrt_flag)
zex_flag \leftarrow ZE \& zx_flag
if( ~vex_flag \& ~zex_flag ) then do
VSR[32xTX+T].dword[0] \leftarrow ConvertSPtoSP64 (result)
VSR[32xTX+T].dword[1] \& 0xUUUU_UUUU_UUUU_UUUU
FPRF }\leftarrow\mathrm{ ClassSP(result)
FR}\leftarrow\textrm{ObU
FI }\leftarrow\textrm{ObU
end
else do
FR}\leftarrow0\textrm{ObO
FI }\leftarrow0\textrm{ObO
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].

A single-precision floating-point estimate of the reciprocal square root of src is placed into doubleword element 0 of \(\operatorname{VSR}[\mathrm{XT}]\) in double-precision format.

Unless the reciprocal of the square root of src would be a zero, an infinity, or a QNaN, the estimate has a relative error in precision no greater than one part in 16384 of the reciprocal of the square root of src. That is,
\[
\left|\frac{\text { estimate }-\frac{1}{\sqrt{s \Gamma c}}}{\frac{1}{\sqrt{\text { SrC }}}}\right| \leq \frac{1}{16384}
\]

Operation with various special values of the operand is summarized below.
\begin{tabular}{ccc}
\hline Source Value & Result & Exception \\
\hline -Infinity & QNaN \(^{1}\) & VXSQRT \\
\hline -Finite & QNaN \(^{1}\) & VXSQRT \\
\hline -Zero & -Infinity \({ }^{2}\) & ZX \\
\hline +Zero & +lnfinity \(^{2}\) & ZX \\
\hline +Infinity & +Zero & None \\
\hline SNaN & QNaN \(^{1}\) & VXSNAN \\
\hline QNaN & QNaN & None \\
\hline
\end{tabular}
1. No result if \(\mathrm{VE}=1\).
2. No result if \(Z E=1\).

The contents of doubleword element 1 of \(\operatorname{VSR}[X T]\) are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to an undefined value. Fl is set to an undefined value.

If a trap-enabled invalid operation exception or a trap-enabled zero divide exception occurs, VSR[XT] and FPRF are not modified.

The results of executing this instruction is permitted to vary between implementations, and between different executions on the same implementation.

\section*{Special Registers Altered}

> FPRF FR=0bU FI=0bU FX OX UX ZX
> XX=0bU VXSNAN VXSQRT

\section*{VSR Data Layout for xsrsqrtesp}
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & unused \\
\hline tgt = VSR[XT] & \\
\hline DP & undefined \\
\hline 0 & 64
\end{tabular}

\section*{VSX Scalar Square Root Double-Precision XX2-form}
\[
\text { xssqrtdp } \quad \mathrm{XT}, \mathrm{XB}
\]
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline -60 & 6 & T & 11 & III & 16 & B & 1 & 75 & BXTX \\
\hline
\end{tabular}
\begin{tabular}{ll}
\(X T\) & \(\leftarrow T X \| T\) \\
\(X B\) & \(\leftarrow B X \| B\)
\end{tabular}
reset_xflags()
\(v\{0:\) inf \(\} \quad \leftarrow\) SquareRootFP(VSR[XB] \{0:63\})
result \(\{0: 63\} \leftarrow \operatorname{RoundToDP}(\) RN, \(v)\)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxsqrt_flag) then SetFX(VXSQRT)
if(xx_flag) then \(\operatorname{SetFX}(X X)\)
vex_flag \(\leftarrow\) VE \& (vxsnan_flag | vxsqrt_flag)
if( ~vex_flag ) then do
VSR[XT] \(\leftarrow\) result || 0xUUUU_UUUU_UUUU_UUUU
FPRF \(\leftarrow\) ClassDP(result)
FR \(\leftarrow\) inc_flag
FI \(\leftarrow\) xx_flag
end
else do
\(\mathrm{FR} \leftarrow 0 \mathrm{~b} 0\)
\(\mathrm{FI} \leftarrow 0 \mathrm{~b} 0\)
end

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].

The unbounded-precision square root of src is produced.

See Table 99.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.
```

Special Registers Altered
FPRF FR FI FX XX VXSNAN VXSQRT

```

VSR Data Layout for xssqrtdp
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|l|l|}
\hline DP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline DP & \multicolumn{2}{c|}{ undefined } \\
\hline 0 & 64 \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \begin{tabular}{l}
\[
\mathrm{v} \leftarrow \mathrm{dQNaN}
\] \\
vxsart_flag \(\leftarrow 1\)
\end{tabular} & \begin{tabular}{l}
\[
\mathrm{V} \leftarrow \mathrm{dQNaN}
\] \\
vxsart_flag \(\leftarrow 1\)
\end{tabular} & \(v \leftarrow+\) Zero & \(v \leftarrow+\) Zero & \(\mathrm{v} \leqslant\) SQRT(src) & \(\mathrm{V} \leftarrow+\) Infinity & V ¢ SrC & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\mathrm{src}) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|ll} 
Explanation: & \\
src & The double-precision floating-point value in doubleword element 0 of VSR[XB]. \\
dQNaN & Default quiet NaN (0x7FF8_0000_0000_0000). \\
NZF & Nonzero finite number. \\
SQRT( x\()\) & The unbounded-precision square root of the floating-point value x. \\
\(\mathrm{Q}(\mathrm{x})\) & Return a QNaN with the payload of x. \\
v & The intermediate result having unbounded signficand precision and unbounded exponent range.
\end{tabular}

Table 99.Actions for xssqrtdp

VSX Scalar Square Root Quad-Precision [using round to Odd] X-form

if MSR. VSX=O then VSX_Unavailable()
reset_xflags()
sic \(\leftarrow\) bfp_CONVERT_FROM_BFP128(VSR[VBB 32 2])
\(v \leftarrow\) bf p_SQUARE_ROOT(sic)
ind \(\leftarrow\) bf P_ROUN_TO_BFP128(RO, FPSCR. RN, v)
result \(\leftarrow\) bf p_CONERT_TO_BFP128(rnd)
if(uxsnan_flag) then SetFX(FPSCR. vXSNaN)
if(uxsort_flag) then SetFX(FPSCR. VXSORT)
if(xx_flag) then Setfx(FPSCR. XX)
\(v x_{-} f l a g \leftarrow v x s n a n\) _flag | vxsqrt_flag
ex_flag \(\leftarrow\) FPSCR.VE \& vx_flag
if ex_flag=0 then do
VSRE[VRT +32\(] \leftarrow\) result
FPSCR.FPRF \(\leftarrow\) forf_Class_bFP128(result)
end
FPSCR. \(F R \leftarrow\left(v x \_f l a g=0\right) \&\) inc_flag
FPSCR. \(F 1 \leftarrow\left(v x_{-}-f \mid \operatorname{lag}=0\right) \& x x_{-}-\overline{1} \mid a g\)
Let src be the floating-point value in VSR[VRB+32] represented in quad-precision format.

If \(\operatorname{src}\) is a Signalling NaN , an Invalid Operation exception occurs and VXSNAN is set to 1 .

If \(\operatorname{src}\) is a negative, non-zero value, an Invalid Operation exception occurs and VXSQRT is set to 1 .

If srC is a Signalling NaN , the result is the Quiet NaN corresponding to src.

Otherwise, if src is a Quiet NaN , the result is src .
Otherwise, if src is a negative value, the result is the default Quiet \(\mathrm{NaN}^{[1]}\).

Otherwise, do the following.
The normalized square root of src is produced with unbounded significand precision and exponent range.

See Table 100, "Actions for xssqrtqp[o]," on page 645.

If \(R O=1\), let the rounding mode be Round to Odd. Otherwise, let the rounding mode be specified by RN. Unless the result is an Infinity or a Zero, the intermediate result is rounded to quad-precision using the specified rounding mode.

See Section 7.3.2.6, "Rounding" on page 383 for a description of rounding modes.

If there is loss of precision, an Inexact exception occurs.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into VSR[ VRT +32 ] in quad-precision format.

FPRF is set to the class and sign of the result. FR is set to indicate if the rounded result was incremented. FI is set to indicate the result is inexact.

If a trap-disabled Invalid Operation exception occurs, \(F P R F\) is set to an undefined value, and FR and FI are set to 0 .

If a trap-enabled Invalid Operation exception occurs, VSR[ VRT +32 ] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

Special Registers Altered:
```

FPRF FR FI FX VXSNAN VXSQRT XX

```

I VSR Data Layout for xssqrtqp[o]
I VSR[VRB+32]
I
I VSR[VRT+32]
I

I 1. The quad-precision default Quiet NaN is the value, \(0 \times 7 F F F_{2} 8000 \_0000 \_0000 \_0000 \_0000 \_0000\).
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \[
\begin{gathered}
v \leftarrow \text { dONaN } \\
\text { vxsqut filag } \leftarrow 1
\end{gathered}
\] & \[
\begin{gathered}
v \leftarrow \operatorname{dONaN} \\
\text { vxspritilag } \leftarrow 1
\end{gathered}
\] & \(v \leftarrow+\) Zer 0 & \(V \leftarrow+\) Ler 0 & \(\mathrm{V} \leqslant \operatorname{sqrit}(\mathrm{src})\) & \(v \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC}\) & \(\mathrm{v} \leftarrow\) quiet(sc) vxsnan filag \(\leftarrow 1\) \\
\hline
\end{tabular}

Explanation:
SIC The quad-precision floating-point value in VSR[ VRB +32\(]\).
dQNaN Default quiet \(\mathrm{NaN}\left(0 \times 7 F F F \_8000_{2} 0000 \_0000 \_0000 \_0000 \_0000\right)\)
NZF Nonzero finite number.
sqrt (x) Return the normalized \({ }^{1}\) square root of floating-point value \(x\), having unbounded significand precision and exponent range.
quiet \((x) \quad\) Convert \(x\) to the corresponding Quiet NaN .
\(v \quad\) The intermediate result having unbounded significand precision and unbounded exponent range.

\section*{\| Table 100. Actions for xssqrtqp[o]}
1. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.

\section*{VSX Scalar Square Root Single-Precision XX2-form}
Xssqrtsp
\begin{tabular}{|l|l|l|l|l|l|l|l|l|}
\hline 60 & & T & T XB \\
\hline 0 & & 6 & & 11 & & 16 & & B \\
\hline
\end{tabular}
```

reset_xflags()
src }\leftarrow\operatorname{VSR[32\timesBX+B].dword[0]
v }\leftarrow\mathrm{ SquareRootDP(src)
result }\leftarrow\mathrm{ RoundToSP(RN,v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxsqrt_flag) then SetFX(VXSQRT)
if(ox_flag) then SetFX(0X)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
vex_flag \leftarrow VE \& (vxsnan_flag | vxsqrt_flag)
if( ~vex_flag ) then do
VSR[32\timesTX+T].dword[0] \leftarrow ConvertToDP(result)
VSR[32xTX+T].dword[1] \& 0xUUUU_UUUU_UUUU_UUUU
FPRF }\leftarrowClasSSP(result
FR}\leftarrow\mathrm{ inc_flag
FI \leftarrowxx_flag
end
else do
FR}\leftarrow0\textrm{b}
FI}\leftarrow\textrm{ObO
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].

The unbounded-precision square root of src is produced.

See Table 99.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT] in double-precision format.

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, \(\mathrm{VSR}[\mathrm{XT}]\) and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

Special Registers Altered
FPRF FR FI FX OX UX XX
VXSNAN VXSQRT
VSR Data Layout for xssqrtsp
src = VSR[XB]

\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|c|}{src} \\
\hline -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \[
\begin{array}{|l|l}
\hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\
\text { vxsqr_flag } \leftarrow 1
\end{array}
\] & \[
\begin{array}{|l|l}
\hline \mathrm{V} \leftarrow \mathrm{dQNaN} \\
\mathrm{vxsart} \text { _lag } \leftarrow 1
\end{array}
\] & \(v \leftarrow+\) Zero & \(v \leftarrow+\) Zero & \(\mathrm{V} \leqslant\) SQRT(src) & \(v \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow \mathrm{SrC}\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src}) \\
& \mathrm{vxsnnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}

\footnotetext{
Explanation:
src The double-precision floating-point value in doubleword element 0 of VSR[XB].
dQNaN Default quiet NaN (0x7FF8_0000_0000_0000).
NZF Nonzero finite number.
\(\operatorname{SQRT}(\mathrm{x}) \quad\) The unbounded-precision and exponent range square root of the floating-point value x .
\(\mathrm{Q}(\mathrm{x}) \quad\) Return a QNaN with the payload of x .
\(v \quad\) The intermediate result having unbounded signficand precision and unbounded exponent range.
}

Table 101.Actions for xssqrtsp

\section*{VSX Scalar Subtract Double-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline xssubdp & & T,XA, & & & \\
\hline \[
\begin{array}{ll} 
& 60 \\
0
\end{array}
\] & \[
\begin{array}{ll} 
& \mathrm{T} \\
6 &
\end{array}
\] & \({ }_{11} \mathrm{~A}\) & 16 & 21 &  \\
\hline \multicolumn{6}{|l|}{XT \(\quad \leftarrow\) TX \| T} \\
\hline \multicolumn{6}{|l|}{\(X A \quad \leftarrow A X \| A\)} \\
\hline \multicolumn{6}{|l|}{} \\
\hline \multicolumn{6}{|l|}{reset_xflags()} \\
\hline \multicolumn{6}{|l|}{src1 \(\leftarrow \operatorname{VSR}[\mathrm{XA}]\{0: 63\}\)} \\
\hline \multicolumn{6}{|l|}{\(\operatorname{src2}\) 2 VSR[XB]\{0:63\}} \\
\hline \multicolumn{6}{|l|}{v\{0:inf\} \(\leftarrow\) AddDP(src1,NegateDP(src2))} \\
\hline \multicolumn{6}{|l|}{result \(\{0: 63\} \leftarrow\) RoundToDP(RN, v)} \\
\hline \multicolumn{6}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{6}{|l|}{if(vxisi_flag) then SetFX(VXISI)} \\
\hline \multicolumn{6}{|l|}{if(ox_flag) then SetFX(0X)} \\
\hline \multicolumn{6}{|l|}{if(ux_flag) then SetFX(UX)} \\
\hline \multicolumn{6}{|l|}{if( \(x\) _f_flag) then \(\operatorname{SetFX}(X X)\)} \\
\hline \multicolumn{6}{|l|}{vex_flag \(\leftarrow\) VE \& (vxsnan_flag | vxisi_flag)} \\
\hline \multicolumn{6}{|l|}{if( ~vex_flag ) then do} \\
\hline \multicolumn{6}{|l|}{VSR[XT] \(\leftarrow\) result || 0xUUUU_UUUU_UUUU_UUUU} \\
\hline \multicolumn{6}{|l|}{FPRF \(\leftarrow\) ClassDP(result)} \\
\hline \multicolumn{6}{|l|}{FR \(\leftarrow\) inc_flag} \\
\hline \multicolumn{6}{|l|}{FI \(\leftarrow x x\) flag} \\
\hline \multicolumn{6}{|l|}{end} \\
\hline \multicolumn{6}{|l|}{else do} \\
\hline FR & \(\leftarrow 0 \mathrm{bo}\) & & & & \\
\hline FI & \(\leftarrow\) 0b0 & & & & \\
\hline end & & & & & \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let src1 be the double-precision floating-point value in doubleword element 0 of VSR[XA].

Let src2 be the double-precision floating-point value in doubleword element 0 of VSR[XB].
src2 is negated and added \({ }^{[1]}\) to src1, producing a sum having unbounded range and precision.

See Table 102.
The sum is normalized \({ }^{[2]}\).
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[XT].

The contents of doubleword element 1 of VSR[XT] are undefined.

FPRF is set to the class and sign of the result. FR is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered}

FPRF FR FI FX OX UX XX
VXSNAN VXISI

VSR Data Layout for xssubdp


\footnotetext{
1. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( \(G, R\), and \(X\) ) enter into the computation.
2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{}} & \multicolumn{8}{|c|}{src2} \\
\hline & & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \multirow{8}{*}{\[
\overline{\mathrm{O}}
\]} & -Infinity & \begin{tabular}{l}
\[
\mathrm{v} \leftarrow \mathrm{dQNaN}
\] \\
vxisi_flag \(\leftarrow 1\)
\end{tabular} & \(\mathrm{v} \leftarrow\)-Infinity & \(\mathrm{v} \leftarrow-\) Infinity & \(v \leftarrow-\) Infinity & \(v \leftarrow-\) Infinity & \(\mathrm{v} \leftarrow-\) Infinity & v \& Src2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { src2 } 2) \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline & -NZF & \(v \leftarrow+\) lnfinity & \(\mathrm{v} \leftarrow \mathrm{S}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{v} \leqslant \mathrm{SrC1}\) & \(\mathrm{v} \leqslant \mathrm{SrC} 1\) & \(\mathrm{v} \leqslant \mathrm{S}(\mathrm{sc} 1, \mathrm{src} 2)\) & \(\mathrm{v} \leftarrow\) - Infinity & v < STC2 & \[
\begin{aligned}
& \hline \begin{array}{l}
v \leftarrow Q(\text { src } 2) \\
v x s n a n \_f l a g \leftarrow 1
\end{array} \\
& \hline
\end{aligned}
\] \\
\hline & -Zero & \(\mathrm{V} \leftarrow+\) lnfinity & \(\mathrm{v} \leftarrow-\mathrm{SrC2}\) & \(v \leftarrow-\) Zero & \(v \leftarrow\) Rezd & \(\mathrm{v} \leftarrow-\mathrm{Sr} \mathrm{C}^{2}\) & \(\mathrm{V} \leftarrow-\) - Infinity & v < Src2 & \[
\begin{array}{|l|}
\hline v \leftarrow Q(s t r c 2) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\] \\
\hline & +Zero & \(\mathrm{v} \leftarrow+\) Infinity & \(\mathrm{v} \leftarrow-\mathrm{Sr} \mathrm{C}^{2}\) & \(v \leqslant\) Rezd & \(v \leftarrow+\) Zero & \(\mathrm{v} \leftarrow-\mathrm{Sr} 2\) & \(\mathrm{V} \leftarrow-\) Infinity & v \& Src2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow Q(\text { srč2) } \\
& \mathrm{vxsnan} \text { _lag } \leftarrow 1
\end{aligned}
\] \\
\hline & +NZF & \(v \leftarrow+\) lnfinity & \(\mathrm{v} \leftarrow \mathrm{S}\) (stc1, src2) & \(\mathrm{v} \leqslant \mathrm{SrC1}\) & \(\mathrm{v} \leqslant\) SrC1 & \(\mathrm{v} \leqslant \mathrm{S}(\mathrm{sc} 1\), src2) & \(v \leftarrow-\) Infinity & v < Scc2 & \[
\begin{aligned}
& \hline \mathrm{v} \leftarrow \mathrm{Q} \text { (src2) } \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline & +Infinity & \(v \leqslant+\) lninity & \(v \leftarrow+\) lnfinity & \(\mathrm{V} \leftarrow+\) Infinity & \(v \leftarrow+\) lnfinity & \(v \leqslant+\) lninity & \(\mathrm{v} \leftarrow \mathrm{dQNaN}\) vxisi_flag \(\leftarrow 1\) & \(\mathrm{v} \leqslant\) Src2 & \[
\begin{aligned}
& v \leftarrow Q(\text { src2 } 2) \\
& v x s n a n \_f l a g \leftarrow 1
\end{aligned}
\] \\
\hline & QNaN & \(\mathrm{v} \leqslant \mathrm{src} 1\) & \(\mathrm{v} \leqslant\) Srci & \(\mathrm{v} \leqslant\) SrC1 & \(\mathrm{v} \leqslant\) Srci & \(\mathrm{v} \leqslant\) Src 1 & \(\mathrm{v} \leqslant\) Src 1 & \(\mathrm{v} \leqslant\) SrCl & \[
\begin{aligned}
& v \leftarrow \text { src1 } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline & SNaN & \(\mathrm{v} \leftarrow \mathrm{Q}\) (scc1) vxsnan_flag \(\leftarrow 1\) & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\text { srcc1) } \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \begin{array}{l}
\mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
& \mathrm{vxsnnan} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|lll}
\hline Explanation: & & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
src1 & The double-precision floating-point value in doubleword element 0 of VSR[XB]. \\
src2 & Default quiet NaN (0x7FF8_0000_0000_0000). \\
dQNaN & Nonzero finite number. \\
NZF & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). \\
Rezd & The floating-point value \(y\) is negated and then added to the floating-point value \(x\). \\
\(\mathrm{S}(\mathrm{x}, \mathrm{y})\) & Return the normalized sum of floating-point value x and negated floating-point value y , having unbounded range and precision. \\
\(\mathrm{S}(\mathrm{x}, \mathrm{y})\) & Note: If \(\mathrm{x}=\mathrm{y}, \mathrm{v}\) is considered to be an exact-zero-difference result (Rezd). \\
\(\mathrm{Q}(\mathrm{x})\) & Return a QNaN with the payload of x.
\end{tabular}

Table 102.Actions for xssubdp

\section*{VSX Scalar Subtract Quad-Precision [using round to Odd] X-form}
\begin{tabular}{lll} 
xssubqp & VRT,VRA,VRB & \((R 0=0)\) \\
xssubqpo & VRT,VRA,VRB & \((R 0=1)\)
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|}
\hline 63 & \({ }_{6}\) VRT & VRA & VRB & & 516 \\
\hline 11
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable\)
reset_xflags()
srcl \& bfp_CONVERT_FROM_ BFP128( VSR[ VRA+32])
src2 \&bfp_CONVERT_FROM_ BFP128( VSR[ VRB+32])
v \&bfp_adD(scci, bfp_NEGate(src2))
rnd \&bfp_ROUND_TO_BFP128(RO, FPSCR, RN, v)
result \& bf p_CONVERT_TO_BFP128(rnd)
if(vxsnan_flag) then SetFX(FPSCR.UXSNAN)
if(uxisi_flag) then SetFX(FPSCR.VXISI)
if(oxfflag) then SetFX(FPSCR.OX)
if(ux_flag) then SetFX(FPSCR.UX)
if(xx_flag) then SetFX(FPSCR. XX)
vx_flag\leftarrowvxsnan_flag| vxisi_flag
ex_flag\leftarrowFPSCR.VE \&vx_flag
if ex_flag=0 then do
VSR[VRT+32] \& result
FPSCR.FPRF \& fprf_CLASS_BFP128(result)
end
FPSCR.FR\&(vx_flag=O) \& inc_flag
FPSCR.FI \& (vx_ flag=0)\& \&x_flag

```

Let srcl be the floating-point value in VSR[VRA+32] represented in quad-precision format.

Let srcl be the floating-point value in VSR[VRB +32 ] represented in quad-precision format.

If either srcl or srcl is a Signalling NaN, an Invalid Operation exception occurs and VXSNAN is set to 1 .

If srcl and \(\mathrm{srcl}^{2}\) are Infinity values having same signs, an Invalid Operation exception occurs and VXISI is set to 1 .

If srcl is a Signalling NaN , the result is the Quiet NaN corresponding to srcl .

Otherwise, if srcl is a Quiet NaN , the result is srcl .
Otherwise, if \(\mathrm{SrCl}_{2}\) is a Signalling NaN , the result is the Quiet NaN corresponding to src .

Otherwise, if src 2 is a Quiet NaN , the result is src .
Otherwise, if srcl and src are Infinity values having same signs, the result is the default Quiet \(\mathrm{NaN}^{[1]}\).

Otherwise, do the following.
The normalized sum of the negation of src 2 added to srcl is produced with unbounded significand precision and exponent range.

See Table 103, "Actions for xssubqp[o]," on page 650.

If the intermediate result is Tiny (i.e., the unbiased exponent is less than -16382) and \(U E=0\), the significand is shifted right \(N\) bits, where \(N\) is the difference between - 16382 and the unbiased exponent of the intermediate result. The exponent of the intermediate result is set to the value - 16382.

If \(R O=1\), let the rounding mode be Round to Odd. Otherwise, let the rounding mode be specified by RN. Unless the result is an Infinity or a Zero, the intermediate result is rounded to quad-precision using the specified rounding mode.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into VSR[ VRT +32 ] in quad-precision format.

FPRF is set to the class and sign of the result. FR is set to indicate if the rounded result was incremented. Fl is set to indicate the result is inexact.

If a trap-disabled Invalid Operation exception occurs, FPRF is set to an undefined value, and FR and FI are set to 0 .

If a trap-enabled Invalid Operation exception occurs, VSR[VRT+32] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered:}
```

FPRF FR FI FX VXSNAN VXISI OX UX XX

```
    VSR Data Layout for xssubqp[o]
VSR[ VRA+32]
\begin{tabular}{l}
\begin{tabular}{|c|}
\hline SrCl \\
\hline VSR[VRB+32] \\
\hline SrC2 \\
\hline VSR[VRT+32] \\
\hline
\end{tabular} \\
\hline
\end{tabular}

Version 3.0
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity &  & & & & & \multirow[t]{5}{*}{\(v \leftarrow \cdot \mid n f i n i t y\)} & \multirow{6}{*}{\(v \leqslant\) Sic2} & \multirow{6}{*}{\[
\begin{aligned}
& v<\text { quiet }(\text { sic } 2) \\
& \text { vxsinan_filag }+1
\end{aligned}
\]} \\
\hline -NZF & \multicolumn{2}{|r|}{\multirow[t]{3}{*}{\(v+\operatorname{sub}(\operatorname{sic}\), sich \()\)
\(v+\operatorname{sich}\)}} & \multicolumn{2}{|c|}{\(v \leqslant\) sicl} & \(v+\operatorname{sub}(\operatorname{sic} 1\), ssc 2\()\) & & & \\
\hline -Zero & & & \(v\) \& Rezd & \(v \leftarrow\) 2ero & \multirow{2}{*}{\(v \leqslant \mathrm{SiCl}\)} & & & \\
\hline --. +Zero & & & \(v \leqslant+\) 2ero & V \& Rezd & & & & \\
\hline の +NZF & & \(v \leftarrow \operatorname{sub}(\operatorname{sic} 1, \mathrm{src} 2)\) & \multicolumn{2}{|c|}{\(v \leftarrow\) sicl} & \(v \leqslant \operatorname{sub}(\) sicl, src 2\()\) & & & \\
\hline +Infinity & \(v \leqslant+\) nfinity & & & & & \[
\begin{gathered}
V \leftarrow \text { OONIN } \\
\text { Vxisi } f \mid \text { ag }<1
\end{gathered}
\] & & \\
\hline QNaN & \multicolumn{5}{|c|}{V + SiCl} & & & \[
\begin{gathered}
v+\text { stcl } \\
\text { vxsinan_ilag }+1 \\
\hline
\end{gathered}
\] \\
\hline SNaN & \multicolumn{8}{|c|}{\[
\begin{aligned}
& \text { v }<\text { quiet (sicl) } \\
& \text { vxsian_ilag }+1
\end{aligned}
\]} \\
\hline \multicolumn{9}{|l|}{Explanation:} \\
\hline sicl & \multicolumn{8}{|l|}{The quad-precision floating-point value in VSR[ VRA 32\(]\).} \\
\hline srcz & \multicolumn{8}{|l|}{The quad-precision floating-point value in VSR[ VRB +32\(]\).} \\
\hline dQNaN & \multicolumn{8}{|l|}{} \\
\hline NZF & \multicolumn{8}{|l|}{Nonzero finite number.} \\
\hline Rezd & \multicolumn{8}{|l|}{Exact-zero-difference result (subtraction of two finite numbers having same magnitude and signs).} \\
\hline sub (x, y) & \multicolumn{8}{|l|}{Return the normalized difference of floating-point value \(x\) and floating-point value \(y\), having unbounded significand precision and exponent range.} \\
\hline & \multicolumn{8}{|l|}{Note: If \(x=y, v\) is considered to be an exact-zero-difference result (Rezd).} \\
\hline quiet ( x ) & \multicolumn{8}{|l|}{Convert x to the corresponding Quiet NaN .} \\
\hline \(v\) & \multicolumn{8}{|l|}{The intermediate result having unbounded significand precision and unbounded exponent range.} \\
\hline
\end{tabular}

\section*{| Table 103. Actions for xssubqp[o]}

\section*{VSX Scalar Subtract Single-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{xssubsp \(\mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)} \\
\hline \[
\begin{array}{ll} 
& 60 \\
0
\end{array}
\] & \[
{ }_{6} \quad \mathrm{~T}
\] & \[
{ }_{11} \mathrm{~A}
\] & \[
{ }_{16} \text { B }
\] & 218 & \begin{tabular}{l}
|ax|BX|TX \\
303031
\end{tabular} \\
\hline
\end{tabular}
reset_xflags()
src1 \(\leftarrow \operatorname{VSR}[32 \times A X+A]\).dword[0]
\(\operatorname{src} 2 \leftarrow \operatorname{VSR}[32 \times B X+B]\).dword \([0]\)
\(\mathrm{v} \quad \leftarrow \operatorname{AddDP}(\operatorname{src} 1, \operatorname{NegateDP}(\operatorname{src} 2))\)
result \(\leftarrow\) RoundToSP (RN, v)
if (vxsnan_flag) then SetFX (VXSNAN)
if (vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(OX)
if (ux_flag) then SetFX(UX)
if (xx_flag) then \(\operatorname{SetFX}(X X)\)
vex_flag \(\leftarrow V E \&\left(v x s n a n \_f l a g \mid v x i s i \_f l a g\right)\)
if ( ~vex_flag ) then do
\(\operatorname{VSR}[32 \times T X+T]\).dword[0] \(\leftarrow\) ConvertSPtoSP64 (result)
VSR [32xTX+T] .dword[1] \(\leftarrow\) 0xUUUU_UUUU_UUUU_UUUU
FPRF \(\leftarrow\) ClassSP (result)
FR \(\leftarrow\) inc_flag
FI \(\leftarrow \mathrm{xx}\) _flag
end
else do
\(\mathrm{FR} \leftarrow 0 \mathrm{bO}\)
\(\mathrm{FI} \leftarrow \mathrm{ObO}\)
end
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let srcl be the double-precision floating-point value in doubleword element 0 of VSR[ XA].

Let \(\mathrm{src}_{2}\) be the double-precision floating-point value in doubleword element 0 of VSR[ XB].
src2 is negated and added \({ }^{[1]}\) to \(\operatorname{srcl}\), producing the sum, v , having unbounded range and precision.

See Table 104, "Actions for xssubsp," on page 652.
\(v\) is normalized \({ }^{[2]}\) and rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element 0 of VSR[ XT].

The contents of doubleword element 1 of VSR[ XT] are undefined.

FPRF is set to the class and sign of the result as represented in single-precision format. \(F R\) is set to indicate if the result was incremented when rounded. Fl is set to indicate the result is inexact.

If a trap-enabled invalid operation exception occurs, VSR[XT] and FPRF are not modified, and FR and FI are set to 0 .

See Table 59, "VSX Scalar Floating-Point Final Result," on page 517.

\section*{Special Registers Altered}
FPRF FR FI FX OX UX XX

VXSNAN VXISI
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{VSR Data Layout for xssubsp
\[
\operatorname{srcl}=\operatorname{VSR}[X A]
\]} \\
\hline DP & unused \\
\hline \multicolumn{2}{|l|}{SrC2 \(=\) VSR[ \(X B]\)} \\
\hline DP & unused \\
\hline \multicolumn{2}{|l|}{tgt \(=\) VSR[ XT\(]\)} \\
\hline DP & undefined \\
\hline 0 & \\
\hline
\end{tabular}

\footnotetext{
1. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( \(G, R\), and \(X\) ) enter into the computation.
2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \[
\begin{aligned}
& \mathrm{V} \leftarrow \text { dONaN } \\
& \text { vxisi_flag } \leftarrow 1
\end{aligned}
\] & \(\mathrm{v} \leftarrow-\) Infinity & \(v \leftarrow-\) Infinity & \(\mathrm{v} \leftarrow-\) Infinity & \(v \leftarrow-\) Infinity & \(\mathrm{v} \leftarrow-\) Infinity & v ¢ STC2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline -NZF & \(\mathrm{v} \leftarrow+\) Infinity & \(\mathrm{v} \leqslant \mathrm{S}\) (stc1, src2) & v ¢ STC1 & v ¢ SCC1 & \(\mathrm{v} \leqslant \mathrm{S}\) (src1, , sc22) & \(\mathrm{V} \leftarrow-\) Infinity & v ¢ STC2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline -Zero & \(\mathrm{V} \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow-\mathrm{SrC2}\) & \(v \leftarrow-\) Zero & \(v \leftarrow\) Rezd & vヶ-SIC2 & \(\mathrm{V} \leftarrow-\) Infinity & v STC2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnn} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline - - +Zero & \(\mathrm{v} \leftarrow+\) Infinity & \(\mathrm{V} \leftarrow-\mathrm{Sr} 22\) & \(v \leftarrow\) Rezd & \(v \leftarrow+\) Zero & \(\mathrm{v} \leftarrow-\mathrm{Sr} \mathrm{C}^{2}\) & \(\mathrm{v} \leftarrow-\) Infinity & v S SC2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} f l a g \leftarrow 1
\end{aligned}
\] \\
\hline 心 +NZF & \(\mathrm{v} \leftarrow+\) Infinity & \(\mathrm{v} \leqslant \mathrm{S}\) (stc1, src2) & v STC1 & v ¢ SCC1 & \(\mathrm{v} \leqslant \mathrm{S}\) (src1, , sc22) & \(\mathrm{v} \leftarrow-\) Infinity & v S SC2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline +Infinity & \(\mathrm{v} \leftarrow+\) Infinity & \(\mathrm{v} \leftarrow+\) Infinity & \(v \leftarrow+\) Infinity & \(v \leftarrow+\) Infinity & \(v \leftarrow+\) Infinity & \begin{tabular}{l}
\[
\mathrm{v} \leftarrow \mathrm{dQNaN}
\] \\
vxisi_flag \(\leftarrow 1\)
\end{tabular} & \(\mathrm{V} \leftarrow\) STC2 & \[
\begin{aligned}
& \hline \begin{array}{l}
\mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
\mathrm{vxsnan} \text { flag } \leftarrow 1
\end{array}
\end{aligned}
\] \\
\hline QNaN & \(v \leftarrow\) SrC1 & v ¢ SCC1 & \(v \leftarrow\) SrC1 & v ¢ SCC1 & v ¢ STC1 & v ¢ SCC1 & \(v \leftarrow\) SrC1 & \[
\begin{aligned}
& \mathrm{v} \leftarrow \text { Src1 } \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline SNaN & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}\left(\mathrm{sr} 1 \mathrm{r}_{1}\right. \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { _lag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { _lag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\text { src1 } 1) \\
& \mathrm{vxsnn} \text { _flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { _lag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\text { src1) } \\
& \mathrm{vxsnn} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{Explanation:} \\
\hline src1 & The double-precision floating-point value in doubleword element 0 of VSR[XA]. \\
\hline src2 & The double-precision floating-point value in doubleword element 0 of VSR[XB]. \\
\hline dQNaN & Default quiet NaN (0x7FF8_0000_0000_0000). \\
\hline NZF & Nonzero finite number. \\
\hline Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). \\
\hline S(x,y) & The floating-point value y is negated and then added to the floating-point value x . \\
\hline S(x,y) & Return the normalized sum of floating-point value \(x\) and negated floating-point value \(y\), having unbounded range and precision. Note: If \(x=y\), \(v\) is considered to be an exact-zero-difference result (Rezd). \\
\hline Q(x) & Return a QNaN with the payload of \(x\). \\
\hline v & The intermediate result having unbounded signficand precision and unbounded exponent range. \\
\hline
\end{tabular}

Table 104.Actions for xssubsp

\section*{VSX Scalar Test for software Divide Double-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{9}{|l|}{xstdivdp BF,XA, XB} \\
\hline 60 & \({ }_{6} \mathrm{BF}\) & // & 11 A & & B & 21 & 61 & \begin{tabular}{|c|c|c|}
\hline\(A X P X\) & 1 \\
29 & 30 & 31 \\
\hline
\end{tabular} \\
\hline XA & \multicolumn{8}{|l|}{\(\leftarrow A X \| A\)} \\
\hline XB & \multicolumn{8}{|l|}{\(\leftarrow \mathrm{BX} \| \mathrm{B}\)} \\
\hline src1 & \multicolumn{8}{|l|}{\(\leftarrow \mathrm{VSR}[\mathrm{XA}]\{0: 63\}\)} \\
\hline src2 & \multicolumn{8}{|l|}{\(\leftarrow \operatorname{VSR}[\mathrm{XB}]\{0: 63\}\)} \\
\hline e_a & \multicolumn{8}{|l|}{\(\leftarrow \operatorname{VSR}[\mathrm{XA}]\{1: 11\}-1023\)} \\
\hline e_b & \multicolumn{8}{|l|}{\(\leftarrow \operatorname{VSR}[\mathrm{XB}]\{1: 11\}-1023\)} \\
\hline fe_flag & \multicolumn{8}{|l|}{\(\leftarrow\) IsNaN(src1) | IsInf(src1) |} \\
\hline & \multicolumn{8}{|c|}{IsNaN(src2) | IsInf(src2) | IsZero(src2) |} \\
\hline & \multicolumn{8}{|c|}{( e_b <= -1022) |} \\
\hline & \multicolumn{8}{|c|}{( e_b >= 1021 ) |} \\
\hline & \multicolumn{8}{|r|}{( ! IsZero(src1) \& ( (e_a - e_b) >= 1023 ) ) |} \\
\hline & \multicolumn{8}{|r|}{( ! IsZero(src1) \& ( (e_a - e_b) <=-1021 ) ) |} \\
\hline & \multicolumn{8}{|c|}{( ! IsZero(src1) \& ( e_a <= -970 ) )} \\
\hline fg_flag & \multicolumn{8}{|l|}{\(\leftarrow \operatorname{IsInf}(\operatorname{src} 1)|\operatorname{IsInf}(\mathrm{src} 2)|\)} \\
\hline & \multicolumn{8}{|c|}{IsZero(src2) | IsDen(src2)} \\
\hline fl_flag & \multicolumn{8}{|l|}{\(\leftarrow\) xsredp_error() <= 2-14} \\
\hline CR[BF] & \multicolumn{8}{|l|}{\(\leftarrow 0 \mathrm{~b} 1\) || fg_flag || fe_flag || 0b0} \\
\hline
\end{tabular}

Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let srcl be the double-precision floating-point value in doubleword element 0 of VSR[ XA].

Let \(\mathrm{src}_{2}\) be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

Let e_a be the unbiased exponent of srcl.
Let \(e_{-} b\) be the unbiased exponent of \(\operatorname{src}\).
\(f e_{\_} f a g\) is set to 1 for any of the following conditions.
- srcl is a NaN or an infinity.
- src2 is a zero, a NaN, or an infinity.
\(-e \quad b\) is less than or equal to -1022.
- e_b is greater than or equal to 1021.
- sicl is not a zero and the difference, e a - e \(b\), is greater than or equal to 1023.
- srcl is not a zero and the difference, e_a - e_b, is less than or equal to -1021.
- srcl is not a zero and \(e_{-}\)a is less than or equal to -970

Otherwise \(f e_{-} f l a g\) is set to 0 .
```

fg_flag is set to 1 for any of the following conditions.
- srcl is an infinity.
- src2 is a zero, an infinity, or a denormalized value.

```

Otherwise \(\mathrm{fg}_{-} f \mathrm{l}\) ag is set to 0 .
\(C R\) field \(B F\) is set to the value Obl ||fg_fag \|fe_fag ||Obo.

\section*{Version 3.0}

VSX Scalar Test for software Square Root Double-Precision XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{xstsqrtdp BF,XB} \\
\hline \[
0
\] & \({ }_{6} \mathrm{BF}\) & & 111 & 16 B & 21 & 106 & \(|\)\begin{tabular}{|c|c|}
\(8 X\) & 1 \\
30 & 31
\end{tabular} \\
\hline XB & \multicolumn{7}{|l|}{\(\leftarrow \mathrm{BX} \| \mathrm{B}\)} \\
\hline src & \multicolumn{7}{|l|}{\(\leftarrow \mathrm{VSR}[\mathrm{XB}]\{0: 63\}\)} \\
\hline e_b & \multicolumn{7}{|l|}{\(\leftarrow \operatorname{VSR}[\mathrm{XB}]\{1: 11\}-1023\)} \\
\hline fe_flag & \multicolumn{7}{|l|}{\(\leftarrow \operatorname{IsNaN}(\mathrm{src})|\operatorname{IsInf}(\mathrm{src})|\) IsZero(src)} \\
\hline \multicolumn{8}{|c|}{IsNeg(src) | ( e_b <=-970)} \\
\hline fg_flag & \multicolumn{7}{|l|}{\(\leftarrow \operatorname{IsInf}(\mathrm{src}) \mid\) IsZero(src) \(\mid\) IsDen(src)} \\
\hline fl_flag & \multicolumn{7}{|l|}{\(\leftarrow\) xsrsqrtedp_error() <= 2-14} \\
\hline CR[BF] & \multicolumn{7}{|l|}{\(\leftarrow 0 b 1\) || fg_flag || fe_flag || 0b0} \\
\hline
\end{tabular}

Let \(X B\) be the value \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[XB].

Let e_b be the unbiased exponent of src.
fe_flag is set to 1 for any of the following conditions.
- src is a zero, a NaN, an infinity, or a negative value.
- e_b is less than or equal to -970

Otherwise fe_flag is set to 0 .
fg_flag is set to 1 for any of the following conditions.
- src is a zero, an infinity, or a denormalized value.

Otherwise fg_flag is set to 0 .
\(C R\) field \(B F\) is set to the value 0b1 || fg_flag || fe_flag || 0b0.

\section*{Special Registers Altered} CR[BF]

\section*{VSR Data Layout for xstsqrtdp}
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|l|ll|}
\hline DP & unused \\
\hline 0 & 64 & 127 \\
\hline
\end{tabular}

VSX Scalar Test Data Class Double-Precision XX2-form
xststdcdp
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 60 & BF \\
\hline 0 & & 6 & DCMX & & & B & \\
\hline
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable()
src \leftarrowVSR[32xBX+B].dword[0]
exponent \& src.bit[1:11]
fraction \leftarrowscc.bit[12:63]
class.Infinity \leftarrow(exponent = Ox>FF) \& (fraction = 0)
class.NaN \&(exponent = Ox7FF)\&(fraction!=0)
class.Zero \leftarrow(exponent = OxOOO) \& (fraction = 0)
class. Denormal }\leftarrow(\mathrm{ exponent = Ox000) \& (fraction!=0)
match
*(DCMX, bit[0] \& class. NaN)

```
CR. bit \([4 \times B F] \leftarrow\) FPSCR. FL \(\leftarrow\) sic.sign
CR. bit \([4 \times B F+1] \leftarrow F P S C R\). FG \(\leftarrow 0 b O\)
CR. bit \(\mid 4 \times 8 F+2] \leftarrow\) FPSCR. FE \(\leftarrow\) match
CR. bit \([4 \times B F+3] \leftarrow F P S C R . F U \leftarrow O b O\)

Let \(X B\) be the sum \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

Bit 0 of \(C R\) field \(B F\) and bit 0 of \(F P C C\) are set to the sign bit of src .

Bit 1 of \(C R\) field \(B F\) and bit 1 of \(F P C C\) are set to \(0 b 0\).
Bit 2 of \(C R\) field \(B F\) and bit 2 of \(F P C C\) are set to indicate whether the data class of src, as represented in double-precision format, matches any of the data classes specified by DCMX (Data Class Mask).
\begin{tabular}{cl} 
DCMX bit & Data Class \\
0 & NaN \\
1 & +Infinity \\
2 & - Infinity \\
3 & +Zero \\
4 & - Zero \\
5 & +Denormal \\
6 & - Denormal
\end{tabular}

Bit 3 of \(C R\) field \(B F\) and bit 3 of \(F P C C\) are set to \(0 b 0\).
Special Registers Altered:
\(C R\) field \(B F\)
FPCC

VSR Data Layout for xststdcdp
SrC \begin{tabular}{ll|l|}
\hline \multicolumn{4}{c|}{ VSR[XB].dword[0] } & unused \\
\hline 0 & 64 & 127 \\
\hline
\end{tabular}

\section*{VSX Scalar Test Data Class Quad-Precision X-form}
xststdcqp
\begin{tabular}{|c|c|c|c|cc|c|}
\hline 63 & BF \\
\hline 0 & & 6 & DCMX & DRB,DCMX \\
\hline
\end{tabular}
if MSR. VSX=O then VSX_Unavailablell
```

STC < USR[VRB+32]
exponent \& src.bit[1:15]

```
fraction \(\leftarrow\) scc.bit[16:127]
class.Infinity \(\leftarrow(\) exponent \(=0 x 7 F F F) \&(\) fraction \(=0)\)
class. NaN \(\leftarrow(\) exponent \(=0 \times 7\) FFF) \& (fraction \(!=0)\)
class.Zero \(\leftarrow(\) exponent \(=0 \times 0000) \&(\) fraction \(=0)\)
class. Denormal \(\leftarrow(\) exponent \(=0 x 0000) \&(\) fraction \(!=0)\)
match \(\quad \leftarrow\) (DCMX.bit[0] \& class. NaN)
        ( DCMX. bit[1] \& class.Infinity \& !sign)
        (DCMX. bit[2] \& class.Infinity \& sign) |
        (DCMX.bit[3] \& class.Zero \& ssign) |
        (DCMX, bit[4] \& class.Zero \& sign) |
        (DCMX, bit[5] \& class. Denormal \& !sign) |
        (DCMX. bitl[6] \& class. Denormal \& sign)
CR.bit[4xBF] \(\leftarrow\) FPSCR.FL \(\leftarrow\) sic.sign
CR. bit \([4 \times B F+1] \leftarrow\) FPSCR. FG \(\leftarrow 060\)
CR. bit \([4 \times B F+2] \leftarrow\) FPSCR. FE \(\leftarrow\) match
CR. bit \([4 \times 8 F+3] \leftarrow\) FPSCR. FU \(\leftarrow 0 b 0\)

Let \(\operatorname{src}\) be the quad-precision floating-point value in VSR[ VRB+32].

Let the DCMX (Data Class Mask) field specify one or more of the 7 possible data classes, where each bit corresponds to a specific data class.
\begin{tabular}{cl} 
DCM bit & Data Class \\
0 & NaN \\
1 & +Infinity \\
2 & - Infinity \\
3 & +Zero \\
4 & - Zero \\
5 & +Denormal \\
6 & - Denormal
\end{tabular}

Bit 0 of \(C R\) field \(B F\) and bit 0 of \(F P C C\) are set to the sign of SrC .

Bit 1 of \(C R\) field \(B F\) and bit 1 of \(F P C C\) are set to \(0 b 0\).
Bit 2 of \(C R\) field \(B F\) and bit 2 of \(F P C C\) are set to indicate whether the data class of sic, as represented in quad-precision format, matches any of the data classes specified by DCM.

Bit 3 of \(C R\) field \(B F\) and bit 3 of \(F P C C\) are set to \(0 b 0\).

\section*{Special Registers Altered:}

CR field \(B F\)
FPCC
| VSR Data Layout for xststdcqp
| VSR[VRB+32]
I sic

VSX Scalar Test Data Class Single-Precision XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{xststdcsp BF,XB,DCMX} \\
\hline \[
60
\] & \({ }_{6} \mathrm{BF}\) & \({ }_{9} \mathrm{DCMX}\) & & & 21 & 298 & | \(3 \times 1 / 18\) \\
\hline
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable(l
src}\leftarrowVSR[32\timesBX+B],dword[0
exponent }\leftarrow\mathrm{ src.bit[1:11]
fraction }\leftarrow\mathrm{ src.bit[12:63]
class.Infinity \leftarrow(exponent = Ox7FF) \& (fraction = 0)
class.NaN }\leftarrow(\mathrm{ exponent = Ox7FF) \& (fraction!= 0)
class.Zero }\leftarrow(\mathrm{ exponent = OxOOO) \& (fraction = 0)
class.Denormal \leftarrow(exponent = OxOOO) \& (fraction ! = 0)
(exponent > OxOOO) \& (exponent < Ox381)
match }\leftarrow(DCMX,bitlO] \& class.NaN
(DCMX,bit[1] \& class.Infinity \& !sign)
(DCMX,bit[2] \& class.Infinity \& sign)
(DCMX,bit[3] \& class.Zero \& !sign)
(DCMX,bit[4] \& class.Zero \& sign)
(DCMX,bit[5] \& class.Denormal \& !sign)
(DCMX,bit[6] \& class.Denormal \& sign)

```
not_SP_value \(\leftarrow(S T C!=\) Convert_SPtoDP(Convert_DPtoSP(src)))
CR, bit [4×BF] \(\leftarrow\) FPSCR.FL \(\leftarrow\) src.sign
CR. bit \([4 \times B F+1] \leftarrow\) FPSCR. FG \(\leftarrow\) ObO
CR. bit \([4 \times B F+2] \leftarrow F P S C R . F E \leftarrow \operatorname{match}\)
\(C R\). bit \([4 \times B F+3] \leftarrow\) FPSCR. FU \(\leftarrow\) not_SP_value

Let \(X B\) be the sum \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

Bit 0 of \(C R\) field \(B F\) and bit 0 of \(F P C C\) are set to the sign bit of src .

Bit 1 of \(C R\) field \(B F\) and bit 1 of \(F P C C\) are set to \(0 b O\).
Bit 2 of \(C R\) field \(B F\) and bit 2 of \(F P C C\) are set to indicate whether the data class of src, as represented in single-precision format, matches any of the data classes specified by DCMX (Data Class Mask).
\begin{tabular}{cl} 
DCMX bit & Data Class \\
0 & NaN \\
1 & +Infinity \\
2 & - Infinity \\
3 & +Zero \\
4 & - Zero \\
5 & +Denormal \\
6 & - Denormal
\end{tabular}

Bit 3 of \(C R\) field \(B F\) and bit 3 of \(F P C C\) are set to indicate if src is not representable in single-precision format.

Special Registers Altered:
CR field BF
FPCC
\| VSR Data Layout for xststdcdp
\(\operatorname{src}\)\begin{tabular}{|l|l|}
\hline & VSR[XB].dword[0] \\
\hline 0 & unused \\
\hline
\end{tabular}

\section*{VSX Scalar Extract Exponent Double-Precision XX2-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{xsxexpdp RT,XB} \\
\hline \[
\begin{array}{|l|l|}
\hline & 60 \\
0 & \\
\hline
\end{array}
\] & \[
{ }_{6} \quad \mathrm{RT}
\] & \[
110
\] & \[
{ }_{16} \mathrm{~B}
\] & \[
347
\] & \(|\)\begin{tabular}{|c|c|}
\(B X\) & 1 \\
30 & 31
\end{tabular} \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailable()
\(\operatorname{src} \leftarrow \operatorname{VSR}[32 \times B X+B]\). dword \([0]\)
GPR[RT] \(\leftarrow(\) SrC \(\gg 52) \& 0 \times 0000 \_0000 \_0000 \_07 F F\)
Let \(X B\) be the sum \(32 \times B X+B\).
Let src be the double-precision floating-point value in doubleword element 0 of VSR[ XB] .

The value of the exponent field in \(\operatorname{src}\) is placed into GPR[RT] in unsigned integer format.

\section*{Special Registers Altered:}

None
Programming Note
This instruction can be used to operate on a single-precision source operand.

VSX Scalar Extract Exponent Quad-Precision X-form
xsxexpqp VRT,VRB
\begin{tabular}{|c|c|c|c|cc|c|}
\hline 63 & \multicolumn{1}{|c|}{ VRT } & \multicolumn{2}{|c|}{2} & VRB & & 804 \\
0 & & & & 11 & & \\
\hline
\end{tabular}
```

if MSR. VSX=0 then VSX_Unavailablel)
sic $\leftarrow$ VSR [ VBB +32$]$
VSR[VRT+32], dword[0] $\leftarrow$ EXTZ64(src, bit [1:15]), 64)
VSR[ VRT+32].dword[1] $\leftarrow$ Ox0000_0000_0000_0000

```

Let src be the quad-precision floating-point value in VSR[ VRB+32].

The contents of the exponent field of \(\operatorname{src}\) (bits 1:15) are zero-extended and placed into doubleword 0 of VSR[ VRT+32].

The contents of doubleword 1 of VSR[ VRT +32] are set to 0.

\section*{Special Registers Altered:}

None

VSR Data Layout for xsxexpdp
\(\square\)

\section*{VSR Data Layout for xsxexpdp}
\begin{tabular}{|c|c|c|c|c|}
\hline src & \multicolumn{4}{|c|}{\(V S R[V R B+32]\)} \\
\hline tgt & VSR[ VRT+32]. dword[0] & & O×0000_0000_0000_0000 & \\
\hline & 0 & \multicolumn{2}{|l|}{63} & 127 \\
\hline
\end{tabular}

\section*{VSX Scalar Extract Significand Double-Precision XX2-form}
Rsxsigdp
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 60 & RT & RT & 1 & & B & & 347 \\
\hline 0 & 6 & & 11 & & 16 & & 21 \\
30 & & 31 \\
\hline
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable()
exponent \leftarrowVSR[32xBX+B], bit[1:11]
fraction \& EXTZ64(VSR[32\timesBX+B]. bit[12:63])
if (exponent != 0) \& (exponent != 2047) then
significand \leftarrowfraction| Ox0010_0000_0000_0000
else
significand \&fraction
GPR[RT] \& significand

```

Let \(X B\) be the sum \(32 \times B X+B\).
Let \(\operatorname{src}\) be the double-precision floating-point value in doubleword element 0 of VSR[ XB].

The significand of \(\operatorname{src}\) is placed into GPR[RT] in unsigned integer format. If src is a normal value, the implicit leading bit is set to 1 .

\section*{Special Registers Altered:}

None

\section*{Programming Note}

This instruction can be used to operate on a single-precision source operand.

\section*{VSX Scalar Extract Significand Quad-Precision X-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline xsxsigq & & VRT,VR & & & & \\
\hline \[
63
\] & \[
{ }_{6} \mathrm{VRT}
\] & \({ }_{11} 18\) & \({ }_{16}\) VRB & 21 & 804 & 11 \\
\hline \multicolumn{7}{|l|}{If MSR.VSX=0 then VSX_Unavailablell} \\
\hline \multicolumn{7}{|l|}{Src \(\leftarrow\) VSR[VRB +32\(]\)} \\
\hline \multicolumn{7}{|l|}{exponent \(\leftarrow\) EXTZ(src.bit[1:15])} \\
\hline \multicolumn{7}{|l|}{fraction \(\leftarrow\) EXTZ128(src.bit[16:127])} \\
\hline \multicolumn{7}{|l|}{If (exponent ! = 0) \& (exponent \(!=32767\) ) then} \\
\hline \multicolumn{7}{|l|}{VSR[VRT+32] ¢fraction \(\mid\) Ox0001_0000_0000_0000_0000_0000_0000_0000} \\
\hline \multicolumn{7}{|l|}{else} \\
\hline \multicolumn{7}{|l|}{VSR[VRT+32] \(\leftarrow\) fraction} \\
\hline
\end{tabular}

Let \(\operatorname{src}\) be the quad-precision floating-point value in VSR[VRB+32].

The significand of src is placed into VSR[ VRT +32].
If the value of the exponent field of \(s r c\) is equal to \(06000 \_0000 \_0000 \_0000\) (i.e., Zero or Denormal value) or Ob111_1111_1111 (i.e., Infinity or NaN), ObO is placed into bit 15 of VSR[VRT+32]. Otherwise (i.e., Normal value), Ob1 is placed into bit 15 of VSR[VRT+32]. The contents of bits \(0: 14\) of VSR[VRT +32 ] are set to 0 .

\section*{Special Registers Altered:}

None

\section*{VSR Data Layout for xsxsigdp}
\(\square\)
VSR Data Layout for xsxsigqp


\section*{Version 3.0}

VSX Vector Absolute Value Double-Precision XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{10}{|l|}{xvabsdp \(\quad\) XT, XB} \\
\hline \[
\begin{array}{ll} 
& 60 \\
0 &
\end{array}
\] & 6 & & & & 16 & B & 21 & 473 & \(\left\lvert\, \begin{aligned} & 3 x X T X \\ & 3031\end{aligned}\right.\) \\
\hline
\end{tabular}
```

XT}\leftarrowTX||
XB}\leftarrowBX|
do i=0 to 127 by 64
VSR[XT]{i:i+63} \leftarrow0b0 || VSR[XB]{i+1:i+63}
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element i from 0 to 1 , do the following. The contents of doubleword element \(i\) of VSR[XB], with bit 0 set to 0 , is placed into doubleword element i of \(\operatorname{VSR}[\mathrm{XT}]\).

Special Registers Altered
None
VSR Data Layout for xvabsdp
src \(=\) VSR[XB]
\begin{tabular}{|c|c|}
\hline DP & DP \\
\hline
\end{tabular}
\[
\operatorname{tgt}=\mathrm{VSR}[\mathrm{XT}]
\]
\begin{tabular}{|l|l|}
\hline DP & DP \\
\hline 0 & 64 \\
\hline
\end{tabular}

VSX Vector Absolute Value Single-Precision XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{xvabssp \(\quad \mathrm{XT}, \mathrm{XB}\)} \\
\hline \[
\begin{array}{ll} 
& 60 \\
0 &
\end{array}
\] & \[
{ }_{6} \mathrm{~T}
\] & \[
{ }_{11} \quad \text { III }
\] & 16 & B & 21 & 409 & BXXTX 3031 \\
\hline
\end{tabular}
```

XT}\leftarrowTX||
XB}\leftarrowBX|
do i=0 to 127 by 32
VSR[XT]{i:i+31}}\leftarrow0b0 || VSR[XB]{i+1:i+31
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element i from 0 to 3 , do the following. The contents of word element \(i\) of \(\operatorname{VSR}[X B]\), with bit 0 set to 0 , is placed into word element i of VSR[XT].

\section*{Special Registers Altered}

None

\section*{VSR Data Layout for xvabssp}
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|l|l|l|l|}
\hline\(S P\) & \(S P\) & \(S P\) & \(S P\) \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|l|l|l|}
\hline SP & \multicolumn{2}{|c|}{SP} & \multicolumn{2}{|c|}{ SP } \\
\hline
\end{tabular}

VSX Vector Add Double-Precision XX3-form
\[
\text { xvadddp } \quad \text { XT,XA,XB }
\]

\begin{tabular}{ll} 
XT & \(\leftarrow \mathrm{TX} \| \mathrm{T}\) \\
XA & \(\leftarrow A X \| A\) \\
XB & \(\leftarrow B X \| B\) \\
ex_flag & \(\leftarrow 0 b 0\)
\end{tabular}
do \(\mathrm{i}=0\) to 127 by 64
reset_xflags()
\(\operatorname{src} 1 \quad \leftarrow \operatorname{VSR}[X A]\{i: i+63\}\)
src2 \(\leftarrow\) VSR[XB]\{i:i \(1+63\}\)
\(\mathrm{v}\{0:\) inf \(\} \quad \leftarrow\) AddDP(src1, src2)
result \(\{\mathrm{i}: 1+63\} \leftarrow \operatorname{RoundToDP}(R N, v)\)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(0x)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
ex_flag \(\leftarrow\) ex_flag | (VE \& vxsnan_flag)
ex_flag \(\leftarrow\) ex_flag | (VE \& vxisi_flag)
ex_flag \(\leftarrow\) ex_flag | ( 0 E \& ox_flag)
ex_flag \(\leftarrow\) ex_flag | (UE \& ux_flag)
ex_flag \(\leftarrow\) ex_flag | (XE \& xx_flag)
end
if( ex_flag \(=0)\) then VSR \([X T] \leftarrow\) result
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element i from 0 to 1 , do the following. Let src1 be the double-precision floating-point operand in doubleword element \(i\) of VSR[XA].

Let src2 be the double-precision floating-point operand in doubleword element \(i\) of VSR[XB].
src2 is added \({ }^{[1]}\) to src1, producing a sum having unbounded range and precision.

The sum is normalized \({ }^{[2]}\).
See Table 105.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element i of VSR[XT] in double-precision format.

See Table 106, "Vector Floating-Point Final Result," on page 663.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

\section*{Special Registers Altered}

FX OX UX XX VXSNAN VXISI
VSR Data Layout for xvadddp
src1 = VSR[XA]
\begin{tabular}{|c|c|}
\hline DP & DP \\
\hline src2 = VSR[XB] & \\
\hline DP & DP \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline DP & \multicolumn{1}{|c|}{ DP } \\
\hline 0 & 64 \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{} & \multicolumn{8}{|c|}{src2} \\
\hline & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline -Infinity & \(v \leftarrow \cdot \| n f i n i t y\) & v ¢-Infinity & \(v \leftarrow \cdot \| n f i n i t y\) & \(v \leftarrow\) Infinity & \(v \leftarrow \cdot\) Infinity & \[
\begin{aligned}
& \hline \begin{array}{l}
\mathrm{v} \leftarrow \mathrm{dQNaN} \\
\text { vxisi_flag } \leftarrow 1
\end{array} \\
& \hline
\end{aligned}
\] & \(\mathrm{V} \leqslant\) Src2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 } 2) \\
& \mathrm{vxsnnn} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline -NZF & \(v \leqslant-\) Infinity & \(\mathrm{v} \leqslant \mathrm{A}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{v} \leqslant \mathrm{SrC} 1\) & \(\mathrm{v} \leqslant\) SrC1 & \(\mathrm{V} \leftarrow A(\) src1, src2) & \(\mathrm{V} \leftarrow+\) Infinity & v < Scc2 & \[
\begin{array}{|l|}
\hline v \leftarrow Q(s r c 2) \\
\text { vxsnan_lag } \leftarrow 1 \\
\hline
\end{array}
\] \\
\hline -Zero & \(v \leftarrow-\) Infinity & v \& Src2 & \(\mathrm{v} \leftarrow\)-Zero & \(v \leftarrow\) Rezd & v ¢ Src2 & \(\mathrm{V} \leftarrow+\) Infinity & v ¢ Src2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\text { src2) } \\
& \mathrm{vxsnn} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline - +Zero & \(v \leftarrow \cdot \| n f i n i t y\) & v \& Src2 & \(v \leftarrow\) Rezd & \(v \leftarrow+\) Zero & v ¢ Sc2 & \(v \leftarrow+\) Infinity & v ¢ Src2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { src2 } 2) \\
& \mathrm{vxsnan} f l \mathrm{lag} \leftarrow 1
\end{aligned}
\] \\
\hline +NZF & \(v \leftarrow\)-nfinity & \(\mathrm{v} \leqslant \mathrm{A}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{v} \leqslant \mathrm{SrC} 1\) & v ¢ SrC1 & \(\mathrm{V} \leftarrow A(\) src1, src 2\()\) & \(\mathrm{v} \leftarrow+\) Infinity & v < Scc2 & \[
\begin{array}{|l|}
\hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 } 2) \\
\mathrm{vxsnan} \text { flag } \leftarrow 1
\end{array}
\] \\
\hline +Infinity & \[
\mathrm{V} \leftarrow \mathrm{dQNaN}
\]
\[
\text { vxisi_flag } \leftarrow 1
\] & \(v \leftarrow+\) Infinity & \(v \leftarrow+\) Infinity & \(v \leftarrow+\) Infinity & \(v \leftarrow+\) lnfinity & \(v \leftarrow+\) Infinity & \(\mathrm{v} \leqslant\) Src2 & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { src2) } \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline QNaN & \(\mathrm{v} \leqslant\) Src 1 & \(\mathrm{v} \leqslant \mathrm{src} 1\) & \(\mathrm{v} \leqslant\) SrC1 & \(v \leftarrow\) src 1 & \(\mathrm{v} \leqslant\) SrC1 & \(\mathrm{v} \leqslant \mathrm{src} 1\) & \(\mathrm{v} \leqslant\) SrC1 & \[
\begin{aligned}
& \mathrm{v} \leftarrow \text { Src1 } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline SNaN & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { src1) } \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \begin{array}{l}
\mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
\mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{array}
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\text { src1) } \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{srcc}) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { _lag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{srccl}) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\text { srcc1) } \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|lll}
\hline Explanation: & \\
src1 & The double-precision floating-point value in doubleword element \(i\) of VSR[XA] (where \(\mathrm{i} \in\{0,1\})\). \\
src 2 & The double-precision floating-point value in doubleword element i of VSR[XB] (where \(\mathrm{i} \in\{0,1\})\). \\
dQNaN & Default quiet NaN (0x7FF8_0000_0000_0000). \\
NZF & Nonzero finite number. \\
Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). \\
\(\mathrm{A}(\mathrm{x}, \mathrm{y})\) & Return the normalized sum of floating-point value x and floating-point value y, having unbounded range and precision. \\
\(\mathrm{Q}(\mathrm{x})\) & Note: If \(\mathrm{x}=-\mathrm{y}, \mathrm{v}\) is considered to be an exact-zero-difference result (Rezd). \\
v & Return a QNaN with the payload of x.
\end{tabular}

Table 105.Actions for xvadddp (element i)
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline Case & ш & Ш & 山 & 山 & \[
\underset{\times}{\boldsymbol{x}}
\] &  &  &  &  & \[
\begin{aligned}
& \text { O } \\
& \underset{T}{4} \\
& 4 \\
& N \\
& N \\
& N \\
& X
\end{aligned}
\] &  & \[
\begin{aligned}
& \text { 『్ } \\
& \text { 4 } \\
& \times \\
& \times
\end{aligned}
\] &  &  &  &  & Returned Results and Status Setting \\
\hline \multirow{17}{*}{Special} & － & － & － & － & － & 0 & 0 & 0 & 0 & 0 & 0 & 0 & － & － & － & － & T（r） \\
\hline & － & － & － & 0 & － & － & － & － & － & － & － & 1 & － & － & － & － & T（r），fx（ZX） \\
\hline & － & － & － & 1 & － & － & － & － & － & － & － & 1 & － & － & － & － & fx（ZX），error（） \\
\hline & 0 & － & － & － & － & － & － & － & － & － & 1 & － & － & － & － & － & T（r），fx（VXSQRT） \\
\hline & 0 & － & － & － & － & － & － & － & － & 1 & － & － & － & － & － & － & T（r），fx（VXZDZ） \\
\hline & 0 & － & － & － & － & － & － & － & 1 & － & － & － & － & － & － & － & T（r），fx（VXIDI） \\
\hline & 0 & － & － & － & － & － & － & 1 & － & － & － & － & － & － & － & － & T（r），fx（VXIISI） \\
\hline & 0 & － & － & － & － & 0 & 1 & － & － & － & － & － & － & － & － & － & T（r），fx（VXIMZ） \\
\hline & 0 & － & － & － & － & 1 & 0 & － & － & － & － & － & － & － & － & － & T（r），fx（VXSNAN） \\
\hline & 0 & － & － & － & － & 1 & 1 & － & － & － & － & － & － & － & － & － & T（r），fx（VXSNAN），fx（VXIMZ） \\
\hline & 1 & － & － & － & － & － & － & － & － & － & 1 & － & － & － & － & － & T（r），fx（VXSQRT） \\
\hline & 1 & － & － & － & － & － & － & － & － & 1 & － & － & － & － & － & － & fx（VXZDZ），error（） \\
\hline & 1 & － & － & － & － & － & － & － & 1 & － & － & － & － & － & － & － & fx（VXIDII），error（） \\
\hline & 1 & － & － & － & － & － & － & 1 & － & － & － & － & － & － & － & － & fx（VXISII），error（） \\
\hline & 1 & － & － & － & － & 0 & 1 & － & － & － & － & － & － & － & － & － & fx（VXIMZ），error（） \\
\hline & 1 & － & － & － & － & 1 & 0 & － & － & － & － & － & － & － & － & － & fx（VXSNAN），error（） \\
\hline & 1 & － & － & － & － & 1 & 1 & － & － & － & － & － & － & － & － & － & fx（VXSNAN），fx（VXIMZ），error） \\
\hline \multirow{5}{*}{Normal} & － & － & － & － & － & － & － & － & － & － & － & － & no & － & － & － & T（r） \\
\hline & － & － & － & － & 0 & － & － & － & － & － & － & － & yes & no & － & － & \(T(r), f(X X)\) \\
\hline & － & － & － & － & 0 & － & － & － & － & － & － & － & yes & yes & － & － & \(\mathrm{T}(\mathrm{r}), \mathrm{fx}(\mathrm{XX})\) \\
\hline & － & － & － & － & 1 & － & － & － & － & － & － & － & yes & no & － & － & \(T(r), f(X X)\) ，error（） \\
\hline & － & － & － & － & 1 & － & － & － & － & － & － & － & yes & yes & － & － & \(T(r), f(X X)\) ，error（） \\
\hline
\end{tabular}

Explanation：
\begin{tabular}{|c|c|}
\hline － & The results do not depend on this condition． \\
\hline \(\mathrm{fx}(\mathrm{x})\) & FX is set to 1 if \(\mathrm{x}=0 . \mathrm{x}\) is set to 1 ． \\
\hline q & The value defined in Table 58，＂Scalar Floating－Point Intermediate Result Handling，＂on page 516，signficand rounded to the target precision，unbounded exponent range． \\
\hline r & The value defined in Table 58，＂Scalar Floating－Point Intermediate Result Handling，＂on page 516，signficand rounded to the target precision，bounded exponent range． \\
\hline \(v\) & The precise intermediate result defined in the instruction having unbounded signficand precision，unbounded exponent range． \\
\hline OX & Floating－Point Overflow Exception status flag，FPSCR \({ }_{\text {Ox }}\) ． \\
\hline error（） & The system error handler is invoked for the trap－enabled exception if the FE0 and FE1 bits in the Machine State Register are set to any mode other than the ignore－exception mode．Update of the target VSR is suppressed for all vector elements． \\
\hline T（x） & The value \(x\) is placed in element \(i\) of VSR［XT］in the target precision format（where \(i \in\{0,1\}\) for results with 64 －bit elements，and \(i\) \(\in\{0,1,3,4\}\) ）for results with 32 －bit elements）． \\
\hline UX & Floating－Point Underflow Exception status flag，FPSCR \({ }_{\text {UX }}\) \\
\hline VXSNAN & Floating－Point Invalid Operation Exception（SNaN）status flag，FPSCR \({ }_{\text {VXSNAN }}\) ． \\
\hline VXSQRT & Floating－Point Invalid Operation Exception（Invalid Square Root）status flag，FPSCR \({ }_{\text {VxSQRT }}\) ． \\
\hline VXIDI & Floating－Point Invalid Operation Exception（Infinity \(\div\) Infinity）status flag，FPSCR \({ }_{\text {VxIDI }}\) ． \\
\hline VXIMZ & Floating－Point Invalid Operation Exception（Infinity \(\times\) Zero）status flag，FPSCR \({ }_{\text {VxIMz }}\) ． \\
\hline VXISI & Floating－Point Invalid Operation Exception（Infinity－Infinity）status flag，FPSCR \({ }_{\text {VxISI }}\)－ \\
\hline VXZDZ & Floating－Point Invalid Operation Exception（Zero \(\div\) Zero）status flag，FPSCR \({ }_{\text {VXZDZ }}\)－ \\
\hline XX & Float－Point Inexact Exception status flag，FPSCR \({ }_{X X}\) ．The flag is a sticky version of FPSCR \(_{F I}\) ．When FPSCR \(_{\text {FI }}\) is set to a new value，the new value of FPSCR \(_{X X}\) is set to the result of ORing the old value of FPSCR \(_{X X}\) with the new value of \(\mathrm{FPSCR}_{\mathrm{FI}}\) ． \\
\hline ZX & Floating－Point Zero Divide Exception status flag，FPSCR \({ }_{\text {Zx }}\) ． \\
\hline
\end{tabular}

Table 106．Vector Floating－Point Final Result

\section*{Version 3.0}


Table 106.Vector Floating-Point Final Result (Continued)

VSX Vector Add Single-Precision XX3-form
xvaddsp
XT, XA, XB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \({ }_{0} 60\) & 6 & T & 11 & A & 16 & B & 21 & 64 & \begin{tabular}{l}
AxBXITX \\
293031
\end{tabular} \\
\hline
\end{tabular}
\begin{tabular}{ll} 
XT & \(\leftarrow \mathrm{TX} \| \mathrm{T}\) \\
XA & \(\leftarrow A X \| A\) \\
XB & \(\leftarrow B X \| B\) \\
ex_flag & \(\leftarrow 0 b 0\)
\end{tabular}
do \(\mathrm{i}=0\) to 127 by 32
reset_xflags()
src1 \(\leftarrow \operatorname{VSR}[X A]\{i: i+31\}\)
src2 \(\leftarrow \operatorname{VSR}[X B]\{i: i+31\}\)
\(\mathrm{v}\{0:\) inf \(\} \quad \leftarrow\) AddSP(src1,src2)
result \(\{\mathrm{i}: 1+31\} \leftarrow \operatorname{RoundTOSP}(R N, v)\)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(0x)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
ex_flag \(\leftarrow\) ex_flag | (VE \& vxsnan_flag)
ex_flag \(\leftarrow\) ex_flag | (VE \& vxisi_flag)
ex_flag \(\leftarrow\) ex_flag | ( 0 E \& ox_flag)
ex_flag \(\leftarrow\) ex_flag \| (UE \& ux_flag)
ex_flag \(\leftarrow\) ex_flag | (XE \& xx_flag)
end
if( ex_flag \(=0)\) then VSR \([X T] \leftarrow\) result
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element i from 0 to 3 , do the following. Let src1 be the single-precision floating-point operand in word element i of VSR[XA].

Let src2 be the single-precision floating-point operand in word element i of VSR[XB].
src2 is added \({ }^{[1]}\) to src1, producing a sum having unbounded range and precision.

The sum is normalized \({ }^{[2]}\).
See Table 107.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into word element i of VSR[XT] in single-precision format.

See Table 106, "Vector Floating-Point Final Result," on page 663.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

\section*{Special Registers Altered}

\section*{FX OX UX XX VXSNAN VXISI}

\section*{VSR Data Layout for xvaddsp}
src1 = VSR[XA]
\begin{tabular}{|c|c|c|c|}
\hline\(S P\) & \(S P\) & \(S P\) & \(S P\) \\
\hline
\end{tabular}
\(\operatorname{src} 2=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|l|l|l|l|}
\hline\(S P\) & \(S P\) & \(S P\) & \(S P\) \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|l|l|l|}
\hline SP & SP & & SP & \multicolumn{2}{|c|}{ SP } \\
\hline 0 & 32 & 64 & 96 & 127 \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{}} & \multicolumn{8}{|c|}{src2} \\
\hline & & -Infinity & -NZF & -Zero & +Zero & +NZF & +Infinity & QNaN & SNaN \\
\hline \multirow{8}{*}{\[
\overline{\mathrm{O}}
\]} & -Infinity & \(\mathrm{v} \leftarrow\) - Infinity & \(\mathrm{v} \leftarrow\) - Infinity & \(v \leftarrow\)-nfinity & \(\mathrm{v} \leftarrow-\operatorname{lnfinity}\) & \(\mathrm{v} \leqslant-\)-nfinity & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \mathrm{vxisis} \text { flag } \leftarrow 1
\end{aligned}
\] & v \& Src2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { _lag } \leftarrow 1
\end{aligned}
\] \\
\hline & -NZF & \(\mathrm{v} \leftarrow\) - Infinity & \(\mathrm{V} \leftarrow \mathrm{A}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{v} \leqslant \mathrm{SrC1}\) & \(\mathrm{v} \leqslant \mathrm{SrC} 1\) & \(\mathrm{v} \leqslant \mathrm{A}(\mathrm{src} 1, \mathrm{src} 2)\) & \(v \leftarrow+\) lnfinity & v < STC2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline & -Zero & \(\mathrm{V} \leftarrow\) - Infinity & v ¢ Sro2 & \(\mathrm{v} \leftarrow\) - Zero & \(v \leftarrow\) Rezd & v ¢ Src2 & \(\mathrm{V} \leftarrow+\) lnfinity & v < Src2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { _lag } \leftarrow 1
\end{aligned}
\] \\
\hline & +Zero & \(\mathrm{v} \leqslant-\)-nfinity & v ¢ Sro2 & \(v \leqslant\) Rezd & \(v \leftarrow+\) Zero & v ¢ SrC2 & \(v \leftarrow+\) Infinity & v \& Src2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { _lag } \leftarrow 1
\end{aligned}
\] \\
\hline & +NZF & v - - Infinity & \(\mathrm{V} \leftarrow \mathrm{A}(\mathrm{src} 1, \mathrm{src} 2)\) & \(\mathrm{v} \leqslant \mathrm{SrC1}\) & \(\mathrm{v} \leqslant \mathrm{SrC} 1\) & \(\mathrm{v} \leqslant \mathrm{A}(\mathrm{src} 1, \mathrm{src} 2)\) & \(v \leftarrow+\) lnfinity & v \& STC2 & \[
\begin{aligned}
& \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] \\
\hline & +Infinity & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{dQNaN} \\
& \text { vxisi_flag } \leftarrow 1
\end{aligned}
\] & \(v \leftarrow+\) lnfinity & \(v \leftarrow+\) lnfinity & \(v \leftarrow+\) Infinity & \(v \leftarrow+\) lnfinity & \(v \leftarrow+\) Infinity & v \& Src2 & \[
\begin{aligned}
& \hline \begin{array}{l}
v \leftarrow Q(s r c 2) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\end{aligned}
\] \\
\hline & QNaN & v ¢ Src1 & \(\mathrm{v} \leqslant\) SrC1 & \(\mathrm{v} \leqslant\) SrC1 & \(\mathrm{v} \leqslant\) SrC1 & v \& SrC1 & \(\mathrm{v} \leqslant\) SrC1 & v \& SrC1 & \[
\begin{aligned}
& v \leftarrow \text { src1 } \\
& \text { vxsnan_flag } \leftarrow 1
\end{aligned}
\] \\
\hline & SNaN & \(\mathrm{v} \leftarrow \mathrm{Q}\) (scc1) vxsnan_flag \(\leftarrow 1\) & \[
\begin{array}{|l|}
\hline v \leftarrow Q(\text { srct1 }) \\
\text { vxsnan_lag } \leftarrow 1
\end{array}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\
& \mathrm{vxsnan} \text { flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
& \mathrm{vxsnan} \text { _flag } \leftarrow 1
\end{aligned}
\] & \[
\begin{aligned}
& \hline \begin{array}{l}
\mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
\text { vxsnan_flag } \leftarrow 1
\end{array}
\end{aligned}
\] & \[
\begin{aligned}
& \mathrm{v} \leftarrow \mathrm{Q}(\text { srccl }) \\
& \mathrm{vxsnnan} \text { _flag } \leftarrow 1
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|ll}
\hline Explanation: & \\
src1 & The single-precision floating-point value in word element \(i\) of VSR[XA] (where \(i \in\{0,1,2,3\})\). \\
src2 & The single-precision floating-point value in word element \(i\) of \(\operatorname{VSR}[X B]\) (where \(i \in\{0,1,2,3\})\). \\
dQNaN & Default quiet NaN (0x7FC0_0000). \\
NZF & Nonzero finite number. \\
Rezd & Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). \\
\(\mathrm{A}(\mathrm{x}, \mathrm{y})\) & Return the normalized sum of floating-point value x and floating-point value y, having unbounded range and precision. \\
\(\mathrm{Q}(\mathrm{x})\) & Note: If \(\mathrm{x}=-\mathrm{y}, \mathrm{v}\) is considered to be an exact-zero-difference result (Rezd). \\
v & Return a QNaN with the payload of x.
\end{tabular}

Table 107.Actions for xvaddsp (element i)

\section*{VSX Vector Compare Equal To Double-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline xvcmpeqd xvcmpeqd & & \[
\begin{aligned}
& \text { XT, XA, } \\
& \text { XT, XA, }
\end{aligned}
\] & \begin{tabular}{l}
(Rc \\
(Rc
\end{tabular} & & & \\
\hline  & \[
\begin{array}{ll} 
& \mathrm{T} \\
6 & \\
\hline
\end{array}
\] & 11 A & 16 B & \(\left|\begin{array}{c}R C \\ 21\end{array}\right| 22\) & 99 & \begin{tabular}{|c|c|c|c|} 
AX & PX \\
29 & 30 & 31 \\
\hline
\end{tabular} \\
\hline \multicolumn{7}{|l|}{XT \(\quad \leftarrow T X \| T\)} \\
\hline \multicolumn{7}{|l|}{\(X A \quad \leftarrow A X \| A\)} \\
\hline \multicolumn{7}{|l|}{\(X B \quad \leftarrow B X \| B\)} \\
\hline \multicolumn{7}{|l|}{ex_flag \(\leftarrow\) 0b0} \\
\hline \multicolumn{7}{|l|}{all_false \(\leftarrow\) 0b1} \\
\hline \multicolumn{7}{|l|}{all_true \(\leftarrow 0 \mathrm{~b} 1\)} \\
\hline \multicolumn{7}{|l|}{do \(i \leftarrow 0\) to 127 by 64 reset_xflags()} \\
\hline \multicolumn{7}{|l|}{src1 \(\leftarrow \operatorname{VSR}[\mathrm{XA}]\{\mathrm{i}: 1+63\}\)} \\
\hline \multicolumn{7}{|l|}{src2 \(\leftarrow\) VSR[XB]\{i: \(i+63\}\)} \\
\hline \multicolumn{7}{|l|}{vxsnan_flag \(\leftarrow\) IsSNaN(src1) | IsSNaN(src2)} \\
\hline if( \(\begin{gathered}\text { Com } \\ \text { resu } \\ \text { all }\end{gathered}\) & pareEQD
lt\{i:i+
false & rcc1, src2
\(\leftarrow 0\) OFF

\(\leftarrow 0000\) & ) then & FFFF & & \\
\hline \multicolumn{7}{|l|}{end} \\
\hline else do
resu
all_ & true & \(\leftarrow 0 \times 00\)
\(\leftarrow\) 0b0 & -0000_00 & 0000 & & \\
\hline end
if(vxsn
ex_flag & an_flag & then SetFx & VXSNAN) & & & \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{if( ex_flag = 0 ) then VSR[XT] \(\leftarrow\) result} \\
\hline \multicolumn{7}{|l|}{if( \(\mathrm{Rc}=1\) ) then do} \\
\hline \multicolumn{7}{|l|}{if( !vex_flag ) then} \\
\hline \multicolumn{7}{|l|}{else} \\
\hline \multicolumn{4}{|l|}{\(\mathrm{CR}[6] \leftarrow\) 0bUUUU} & & & \\
\hline
\end{tabular}

Two zero inputs of same or different signs return true for that element.

Two infinity inputs of same signs return true for that element.

If \(R c=1, C R\) Field 6 is set as follows.
- Bit 0 of \(C R[6]\) is set to indicate all vector elements compared true.
- Bit 1 of CR[6] is set to 0 .
- Bit 2 of \(\operatorname{CR}[6]\) is set to indicate all vector elements compared false.
- Bit 3 of \(C R[6]\) is set to 0 .

If a trap-enabled exception occurs in any element of the vector, no results are written to \(\operatorname{VSR}[\mathrm{XT}]\) and the contents of \(\mathrm{CR}[6]\) are undefined if Rc is equal to 1 .

\section*{Special Registers Altered}
```

    CR[6] . . . . . . . . . . . . . . . . . . . . . . . . (if Rc=1)
    FX VXSNAN
    ```

\section*{VSR Data Layout for xvempeqdp[.]}
src1 = VSR[XA]
\begin{tabular}{|c|c|}
\hline DP & DP \\
\hline src2 = VSR[XB] & \\
\hline DP & DP \\
\hline
\end{tabular}
\(\operatorname{tgt}=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline & \multicolumn{1}{|c|}{ MD } \\
\hline 0 & 64 \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element i from 0 to 1 , do the following. Let src1 be the double-precision floating-point operand in doubleword element \(i\) of VSR[XA].

Let src2 be the double-precision floating-point operand in doubleword element \(i\) of VSR[XB].
src1 is compared to src2.
The contents of doubleword element i of VSR[XT] are set to all 1 s if src1 is equal to \(\operatorname{src} 2\), and is set to all Os otherwise.

A NaN input causes the comparison to return false for that element.

\section*{VSX Vector Compare Equal To} Single-Precision XX3-form
\begin{tabular}{lll} 
xvcmpeqsp & \(X T, X A, X B\) & \((\mathrm{Rc}=0)\) \\
xvcmpeqsp. & \(\mathrm{XT}, \mathrm{XA}, \mathrm{XB}\) & \((\mathrm{Rc}=1)\)
\end{tabular}


do \(\mathrm{i}=0\) to 127 by 32
    reset_xflags()
    \(\operatorname{src} 1 \quad \leftarrow \operatorname{VSR}[X A]\{i: i+31\}\)
    src2 \(\leftarrow\) VSR[XB] \(\{i: i+31\}\)
    vxsnan_flag \(\leftarrow\) IsSNaN(src1) \| IsSNaN(src2)
    f( CompareEQSP(src1, src2) ) then
        result \(\{i: i+31\} \leftarrow 0 x F F F F \_F F F F\)
    end
    se do
        result \(\{i: i+31\} \leftarrow 0 x 0000 \_0000\)
        all_true \(\leftarrow 0 \mathrm{bb} 0\)
    end
    if(vxsnan_flag) then SetFX(VXSNAN)
    ex_flag \(\leftarrow\) ex_flag \| (VE \& vxsnan_flag)
if \((\) ex_flag \(=0)\) then VSR[XT] \(\leftarrow\) result
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 x A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

For each vector element i from 0 to 3 , do the following. Let src1 be the single-precision floating-point operand in word element i of VSR[XA].

Let src2 be the single-precision floating-point operand in word element i of VSR[XB].
src1 is compared to src2.
The contents of word element i of VSR[XT] are set to all 1 s if src1 is equal to src 2 , and is set to all Os otherwise.

A NaN input causes the comparison to return false for that element.

Two zero inputs of same or different signs return true for that element.

Two infinity inputs of same signs return true for that element.

If \(R \mathrm{c}=1, \mathrm{CR}\) Field 6 is set as follows.
- Bit 0 of \(C R[6]\) is set to indicate all vector elements compared true.
- Bit 1 of CR[6] is set to 0 .
- Bit 2 of CR[6] is set to indicate all vector elements compared false.
- Bit 3 of CR[6] is set to 0 .

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT] and the contents of \(C R[6]\) are undefined if Rc is equal to 1 .

\section*{Special Registers Altered}
```

CR[6]
(if $\mathrm{Rc}=1$ )
FX VXSNAN

```

\section*{VSR Data Layout for xvcmpeqsp[.]}
src1 = VSR[XA]
\begin{tabular}{|c|c|c|c|}
\hline SP & SP & SP & SP \\
\hline Src2 \(=V\) VR[XB] \\
\hline SP & \(S P\) & \(S P\) & \(S P\) \\
\hline
\end{tabular}
tgt \(=\) VSR \([\mathrm{XT}]\)
\begin{tabular}{|l|l|l|l|}
\hline MW & MW & \multicolumn{2}{|c|}{ MW } \\
\hline
\end{tabular}

\section*{VSX Vector Compare Greater Than or Equal To Double-Precision XX3-form}
```

xvcmpgedp XT,XA,XB (Rc=0)
xvcmpgedp. XT,XA,XB (Rc=1)

```

```

| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |
| all_false | $\leftarrow 0 b 1$ |
| all_true | $\leftarrow 0 b 1$ |
|  |  |
| do i=0 to 127 by 64 |  |
| $\quad$ reset_xflags () |  |
|  | src1 |
|  | $\leftarrow \operatorname{VSR}[X A]\{i: i+63\}$ |
| src2 | $\leftarrow \operatorname{VSR}[X B]\{i: i+63\}$ |

    if( IsSNaN(src1) | IsSNaN(src2) ) then do
            vxsnan_flag \leftarrow 0b1
            if(VE=0) then vxvc_flag \leftarrow0b1
        end
        else vxvc_flag \leftarrow IsQNaN(src1) | IsQNaN(src2)
        if( CompareGEDP(src1,src2) ) then
            result{i:i+63} \leftarrow0xFFFF_FFFF_FFFF_FFFF
            all_false }\leftarrow00
        end
        else do
                result{i:i+63} \leftarrow 0x0000_0000_0000_0000
            all_true }\leftarrow0\mathrm{ 0b0
    end
    if(vxsnan_flag) then SetFX(VXSNAN)
    if(vxvc_flag) then SetFX(VXVC)
    ex_flag \leftarrow ex_flag | (VE & vxsnan_flag)
        ex_flag \leftarrow ex_flag | (VE & vxvc_flag)
    end
if( ex_flag = 0 ) then VSR[XT] \leftarrow result
if(Rc=1) then do
if( !vex_flag ) then
CR[6] \leftarrow all_true || 0b0 || all_false || 0b0
else
CR[6] \leftarrow 0bUUUU
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element i from 0 to 1, do the following. Let src1 be the double-precision floating-point operand in doubleword element \(i\) of VSR[XA].

Let src2 be the double-precision floating-point operand in doubleword element \(i\) of VSR[XB].
src1 is compared to src2.

VSX Vector Compare Greater Than or Equal To Single-Precision XX3-form


The contents of word element \(i\) of VSR[XT] are set to all 1s if src1 is greater than or equal to src2, and is set to all Os otherwise.

A NaN input causes the comparison to return false for that element.

Two zero inputs of same or different signs return true for that element.

Two infinity inputs of same signs return true for that element.

If \(\mathrm{Rc}=1, \mathrm{CR}\) Field 6 is set as follows.
- Bit 0 of \(C R[6]\) is set to indicate all vector elements compared true.
- Bit 1 of \(\mathrm{CR}[6]\) is set to 0 .
- Bit 2 of \(C R[6]\) is set to indicate all vector elements compared false.
- Bit 3 of \(C R[6]\) is set to 0 .

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT] and the contents of \(C R[6]\) are undefined if Rc is equal to 1 .

\section*{Special Registers Altered}
```

CR[6]
(if $\mathrm{Rc}=1$ )
FX VXSNAN VXVC

```

VSR Data Layout for xvcmpgesp[.]
src1 \(=\) VSR[XA]
\begin{tabular}{|c|c|c|c|}
\hline\(S P\) & \(S P\) & \(S P\) & \(S P\) \\
\hline Src2 \(=V\) VR[XB] \\
\hline SP & \(S P\) & \(S P\) & \(S P\) \\
\hline
\end{tabular}
tgt = VSR[XT]
\begin{tabular}{|l|l|l|l|}
\hline MW & \multicolumn{2}{|c|}{ MW } & MW \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 x A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element i from 0 to 3 , do the following.
Let src1 be the single-precision floating-point operand in word element \(i\) of \(\operatorname{VSR}[X A]\).

Let src2 be the single-precision floating-point operand in word element \(i\) of VSR[XB].
src1 is compared to src2.

\section*{VSX Vector Compare Greater Than} Double-Precision XX3-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline xvcmpgtd xvcmpgtd & & \[
\begin{aligned}
& \text { KT,XA, } \\
& \text { k } \mathrm{T}, \mathrm{XA},
\end{aligned}
\] & \[
\begin{aligned}
& (\mathrm{Rc}= \\
& (\mathrm{Rc}=
\end{aligned}
\] & & & \\
\hline \[
\begin{array}{ll} 
& 60 \\
0
\end{array}
\] & \(6_{6}{ }^{\text {T }}\) & \({ }_{11} \mathrm{~A}\) & 16 & & 107 &  \\
\hline \multicolumn{7}{|l|}{XT \(\quad \leftarrow \mathrm{TX} \|\) T} \\
\hline \multicolumn{7}{|l|}{\(X A \quad \leftarrow A X \| A\)} \\
\hline \multicolumn{7}{|l|}{\(X B \quad \leftarrow B X \| B\)} \\
\hline \multicolumn{7}{|l|}{ex_flag \(\leftarrow 0 \mathrm{bb}\)} \\
\hline \multicolumn{7}{|l|}{all_false \(\leftarrow 0 \mathrm{~b} 1\)} \\
\hline \multicolumn{7}{|l|}{all_true \(\leftarrow 0 \mathrm{~b} 1\)} \\
\hline \multicolumn{7}{|l|}{do i=0 to 127 by 64} \\
\hline \multicolumn{7}{|l|}{reset_xflags()} \\
\hline \multicolumn{7}{|l|}{src1 \(\leftarrow\) VSR[XA]\{i:i+63\}} \\
\hline \multicolumn{7}{|l|}{src2 \(\leftarrow\) VSR[XB]\{i:i+63\}} \\
\hline \multicolumn{7}{|l|}{if( IsSNaN(src1) | IsSNaN(src2) ) then do} \\
\hline \multicolumn{7}{|c|}{vxsnan_flag \(\leftarrow 0 \mathrm{~b} 1\)} \\
\hline \multicolumn{7}{|c|}{if \((\mathrm{VE}=0)\) then vxvc _flag \(\leftarrow 0 \mathrm{~b} 1\)} \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{else vxvc_flag \(\leftarrow\) IsQNaN(src1) | IsQNaN(src2)} \\
\hline \multicolumn{7}{|l|}{if( CompareGTDP(src1,src2) ) then do} \\
\hline \multicolumn{7}{|c|}{result \(\{1: i+63\} \leftarrow 0 x F F F F \_F F F F \_F F F F\) FFFFF} \\
\hline \multicolumn{7}{|c|}{all_false \(\leftarrow\) 0b0} \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{else do} \\
\hline \multicolumn{7}{|c|}{result \(\{i: i+63\} \leftarrow\) 0x0000_0000_0000_0000} \\
\hline \multicolumn{7}{|c|}{all_true \(\leftarrow\) 0b0} \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{7}{|l|}{if(vxvc_flag) then SetFX(VXVC)} \\
\hline \multicolumn{7}{|l|}{ex_flag \(\leftarrow\) ex_flag \| (VE \& vxsnan_flag)} \\
\hline \multicolumn{7}{|l|}{ex_flag \(\leftarrow\) ex_flag | (VE \& vxvc_flag)} \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{if( ex_flag = 0 ) then VSR[XT] \(\leftarrow\) result} \\
\hline \multicolumn{7}{|l|}{if( \(\mathrm{Rc}=1\) ) then do} \\
\hline \multicolumn{7}{|l|}{if( !vex_flag ) then} \\
\hline \multicolumn{7}{|c|}{CR[6] \(\leftarrow\) all_true || 0b0 || all_false || 0b0} \\
\hline \multicolumn{7}{|l|}{else} \\
\hline \multicolumn{7}{|c|}{\(C R[6] \leftarrow\) ObUUUU} \\
\hline end & & & & & & \\
\hline
\end{tabular}

The contents of doubleword element i of VSR[XT] are set to all 1 s if src1 is greater than src2, and is set to all 0 s otherwise.

A NaN input causes the comparison to return false for that element.

Two zero inputs of same or different signs return false for that element.

If \(R c=1, C R\) Field 6 is set as follows.
- Bit 0 of \(C R[6]\) is set to indicate all vector elements compared true.
- Bit 1 of \(C R[6]\) is set to 0 .
- Bit 2 of \(\operatorname{CR}[6]\) is set to indicate all vector elements compared false.
- Bit 3 of \(C R[6]\) is set to 0 .

If a trap-enabled exception occurs in any element of the vector, no results are written to \(\operatorname{VSR}[\mathrm{XT}]\) and the contents of \(C R[6]\) are undefined if \(R c\) is equal to 1.

\section*{Special Registers Altered}
\[
\begin{aligned}
& \text { CR[6] ........................................ } \mathrm{Rc}=1 \text { ) } \\
& \text { FX VXSNAN VXVC }
\end{aligned}
\]

VSR Data Layout for xvcmpgtdp[.]
src1 = VSR[XA]
\begin{tabular}{|c|c|}
\hline DP & DP \\
\hline
\end{tabular}
src2 = VSR[XB]
\begin{tabular}{|c|c|}
\hline DP & DP \\
\hline
\end{tabular}
tgt \(=\) VSR \([\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline & \multicolumn{2}{|c|}{ MD } \\
\hline 0 & 64 \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element \(i\) from 0 to 1 , do the following. Let srcl be the double-precision floating-point operand in doubleword element \(i\) of VSR[XA].

Let srcl be the double-precision floating-point operand in doubleword element \(i\) of VSR[ XB].
srcl is compared to srcl .

VSX Vector Compare Greater Than Single-Precision XX3-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline xvcmpgts xvcmpgtsp & & \[
\begin{aligned}
& \text { KT,XA, } \\
& \text { KT,XA, }
\end{aligned}
\] & \begin{tabular}{l}
(Rc \\
(Rc
\end{tabular} & & & \\
\hline \[
\begin{array}{ll} 
& 60 \\
0
\end{array}
\] & \[
{ }_{6} \quad \mathrm{~T}
\] & \({ }_{11} \mathrm{~A}\) & 16 B & & & \begin{tabular}{|c|c|c|} 
AX & BX \\
29 & 30 \\
30 & 31 \\
\hline
\end{tabular} \\
\hline \multicolumn{7}{|l|}{XT \(\quad \leftarrow\) TX \| T} \\
\hline \multicolumn{7}{|l|}{\(X A \quad \leftarrow A X \| A\)} \\
\hline \multicolumn{7}{|l|}{\(X B \quad \leftarrow \mathrm{BX} \| \mathrm{B}\)} \\
\hline \multicolumn{7}{|l|}{ex_flag \(\leftarrow 0 \mathrm{~b} 0\)} \\
\hline \multicolumn{7}{|l|}{all_false \(\leftarrow 0 \mathrm{~b} 1\)} \\
\hline \multicolumn{7}{|l|}{all_true \(\leftarrow 0 \mathrm{~b} 1\)} \\
\hline \multicolumn{7}{|l|}{do i=0 to 127 by 32} \\
\hline \multicolumn{7}{|l|}{reset_xflags()} \\
\hline \multicolumn{7}{|l|}{src1 \(\leftarrow\) VSR[XA]\{i:i+31\}} \\
\hline \multicolumn{7}{|l|}{\(\operatorname{src2}\) 2 \(\leftarrow \mathrm{VSR}[\mathrm{XB}]\{i: i+31\}\)} \\
\hline \multicolumn{7}{|l|}{if( IsSNaN(src1) | IsSNaN(src2) ) then do} \\
\hline \multicolumn{7}{|l|}{vxsnan_flag \(\leftarrow 0 \mathrm{~b} 1\)} \\
\hline \multicolumn{7}{|c|}{if \((\mathrm{VE}=0)\) then vxvc_flag \(\leftarrow 0 \mathrm{~b} 1\)} \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{else vxvc_flag \(\leftarrow\) IsQNaN(src1) | IsQNaN(src2)} \\
\hline \[
\begin{aligned}
& \text { if( Com } \\
& \text { resu } \\
& \text { all }
\end{aligned}
\] & pareGTS
lt \(\{1: 1+\)
false & rc1, src 2
\(\leftarrow 0 \mathrm{XFF}\)
\(\leftarrow 60 \mathrm{~b} 0\) & ) then & & & \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{else do} \\
\hline \multicolumn{7}{|c|}{result \(\{i: 1+31\} \leftarrow 0 x 0000 \_0000\)} \\
\hline \multicolumn{7}{|c|}{all_true \(\leftarrow 00 \mathrm{~b} 0\)} \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{if(vxsnan_flag) then SetFX(VXSNAN)} \\
\hline \multicolumn{7}{|l|}{if(vxvc_flag) then SetFX(VXVC)} \\
\hline \multicolumn{7}{|l|}{ex_flag \(\leftarrow\) ex_flag | (VE \& vxsnan_flag)} \\
\hline \multicolumn{7}{|l|}{ex_flag \(\leftarrow\) ex_flag | (VE \& vxvc_flag)} \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{if( ex_flag = 0 ) then VSR[XT] \(\leftarrow\) result} \\
\hline \multicolumn{7}{|l|}{if( \(\mathrm{Rc}=1)\) then do} \\
\hline \multicolumn{7}{|l|}{if( !vex_flag ) then} \\
\hline \multicolumn{7}{|l|}{CR[6] \(\leftarrow\) all_true || 0b0 || all_false || 0b0} \\
\hline \multicolumn{7}{|l|}{else} \\
\hline \multicolumn{7}{|c|}{\(\mathrm{CR}[6] \leftarrow\) 0bUUUU} \\
\hline end & & & & & & \\
\hline
\end{tabular}

The contents of word element \(i\) of VSR[ XT] are set to all 1 s if srcl is greater than srcc , and is set to all Os otherwise.

A NaN input causes the comparison to return false for that element.

Two zero inputs of same or different signs return false for that element.

If \(R C=1, C R\) Field 6 is set as follows.
- Bit 0 of \(C R[6]\) is set to indicate all vector elements compared true.
- Bit 1 of \(C R[6]\) is set to 0 .
- Bit 2 of CR[6] is set to indicate all vector elements compared false.
- Bit 3 of \(C R[6]\) is set to 0 .

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT] and the contents of \(C R[6]\) are undefined if \(R C\) is equal to 1 .

\section*{Special Registers Altered}
```

CR[6]
(if Rc=1)
FX VXSNAN VXVC

```

\section*{VSR Data Layout for xvempgtsp[.]}
src1 = VSR[XA]
\begin{tabular}{|l|l|l|l|}
\hline\(S P\) & SP & SP & SP \\
\hline
\end{tabular}
\(\operatorname{src} 2=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|c|c|}
\hline\(S P\) & \(S P\) & \(S P\) & \(S P\) \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|l|l|}
\hline MW & MW & \multicolumn{2}{|c|}{ MW } \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element ifrom 0 to 3 , do the following.
Let srcl be the single-precision floating-point operand in word element \(i\) of VSR[XA].

Let srcl be the single-precision floating-point operand in word element \(i\) of VSR[XB].
srcl is compared to src 2 .

\section*{VSX Vector Compare Not Equal Double-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{2}{*}{xvcmpnedp xvcmpnedp.}} & \multicolumn{4}{|l|}{\multirow[t]{2}{*}{VRT,VRA,VRB
VRT,VRA,VRB}} & & \multirow[t]{2}{*}{\[
\begin{aligned}
& (R C=0) \\
& (R C=1)
\end{aligned}
\]} \\
\hline & & & & & & & \\
\hline \[
60
\] & \[
{ }_{6} \mathrm{~T}
\] & \[
{ }_{11} \quad \mathrm{~A}
\] & 16 & B & \[
\left\lvert\, \begin{aligned}
& \mathrm{RC} \\
& 21
\end{aligned} \underbrace{}_{22}\right.
\] & 123 & \[
\left|\begin{array}{l}
2 x|x| x+x|x| \\
2930 \mid 31
\end{array}\right|
\] \\
\hline
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable()
all true }
all_false \leftarrow1
reset_flags()
do i = 0 to l
srcl \leftarrow bfp_CONVERT_FROM_BFP64(VSR[32xAX+A],dword[i])
sic2 \leftarrow bfp_CONVERT_FROM_BFP64(VSR[32\timesBX+B],dword[i])
vxsnan_flag \leftarrowvxsnan_flag | (srcl.type="SNaN")
| (src2.type="SNaN")
if bfp_COMPARE_NE(srcl, src2)=1 then do
result.dwordi[0] <OxFFFF_FFFF_FFFF_FFFF
all_false}\leftarrow
end
else do
result.dword[0]}\leftarrow0\times0000_0000_0000_000
all_true}\leftarrow
end
end
vex_flag \leftarrowFPSCR.VE\&vxsnan_flag
if (uxsnan_flag=1) SetFX(FPSCR.vXSNAN)
if (vex_flag=0) then VSR[32xTX+T]}\leftarrow\mathrm{ result
if Rc=1 then do
CR.bit[56] \leftarrowall_true
CR.bit[57] }\leftarrow0b
CR.bit[58]}\leftarrowall_fals
CR.bit[59]}\leftarrow\textrm{ObO
end

```

\section*{VSX Vector Compare Not Equal Single-Precision XX3-form}
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{2}{|l|}{xvcmpnesp} & \multicolumn{2}{|l|}{\multirow[t]{2}{*}{\begin{tabular}{l}
VRT,VRA,VRB \\
VRT,VRA,VRB
\end{tabular}}} & & & ( \(R C=0\) ) \\
\hline xvcmpn & & & & & & ( \(R C=1\) ) \\
\hline \[
\begin{array}{ll} 
& 60 \\
0
\end{array}
\] & \[
6 \quad \text { T }
\] & \({ }_{11} \mathrm{~A}\) & \({ }_{16}\) B & \(|\)\begin{tabular}{|c|}
\(R c\) \\
21 \\
22
\end{tabular} & 91 &  \\
\hline
\end{tabular}

The contents of doubleword 1 of VSR[VRT] are set to \(0 \times 0000 \_0000 \_0000 \_0000\).

If a trap-enabled Invalid Operation occurs, VSR[XT] is not modified.

Special Registers Altered:
FX VXSNAN
```

if MSR.VSX=0 then VSX_Unavailablell
all_true }\leftarrow
all_false \leftarrow }\leftarrow
reset_flags()
do i = 0 to 3
srcl \leftarrow bfp_CONVERT_FROM_BFP32(VSR[32xAX+A], word[i])
srC2 \leftarrowbfp_CONVERT_FROM_BFP32(VSR[32\timesBX+B], word[i])
vxsnan_flag \leftarrowvxsnan_flag ( (srcl.type="SNaN")
| (src2.type="SNaN")
if bfp_COMPARE NE(srcl, src2)=1 then do
result.word[0] < OXFFFF_FFFF
all_false}\leftarrow
end
else do
result.word[0]}\leftarrow0\times0000_000
all_true}\leftarrow
end
end
vex_flag \leftarrowFPSCR.VE\& vxsnan_flag
if (vxsnan_flag=1) SetFX(FPSCR.VXSNAN)
if (vex_flag=0) then VSR[32xTX+T] \leftarrow result
if RC=1 then do
CR.bit[56]}\leftarrow\mathrm{ all_true
CR.bit[57]}\leftarrow\mp@subsup{\textrm{ObO}}{}{-
CR.bit[58]}\leftarrow\mathrm{ all_false
CR.bit[59]}\leftarrow0b
end

```
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

Let srcl be the double-precision floating-point value in doubleword 0 of VSR[ XA].

Let \(\operatorname{src2}\) be the double-precision floating-point value in doubleword 0 of VSR[ XB].
srcl is compared to src .
A NaN compared to any value, including itself, compares true for the predicate, not equal.

The contents of doubleword 0 of VSR[VRT] are set to OXFFFF_FFFFFFFF_FFFF if \(\mathrm{SrCl}_{1}\) is greater than src 2 , and are set to \(0 \times 0000 \_0000 \_0000 \_0000\) otherwise.

VSX Vector Copy Sign Double-Precision XX3-form
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline xvcpsg & & XT, XA & & & & & \\
\hline \({ }^{6} 60\) & \({ }_{6}{ }^{\text {T }}\) & \({ }_{11}\) & & & 21 & 240 &  \\
\hline
\end{tabular}
```

XT}\leftarrowTX||
XA\leftarrowAX|A
XB}\leftarrowBX|
do i=0 to 127 by 64
VSR[XT]{i:i+63} \leftarrow VSR[XA]{i} || VSR[XB]{i+1:i+63}
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element i from 0 to 1 , do the following.
The contents of bit 0 of doubleword element i of \(\mathrm{VSR}[\mathrm{XA}]\) are concatenated with the contents of bits 1:63 of doubleword element \(i\) of \(\operatorname{VSR}[X B]\) and placed into doubleword element \(i\) of \(\operatorname{VSR}[\mathrm{XT}]\).

\section*{Special Registers Altered}

None
\begin{tabular}{ll}
\hline Extended Mnemonic & Equivalent To \\
\hline xvmovdp \(\quad \mathrm{XT}, \mathrm{XB}\) & \(\mathrm{xvcpsgndp} \mathrm{XT}, \mathrm{XB}, \mathrm{XB}\) \\
\hline
\end{tabular}

Table 108:


VSX Vector Copy Sign Single-Precision XX3-form
xvcpsgnsp XT,XA,XB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & 11 & A & 16 & B & 21 & 208 &  \\
\hline
\end{tabular}
```

XT}\leftarrowTX||
XA\leftarrowAX|A
XB}\leftarrowBX|
do i=0 to 127 by 32
VSR[XT]{i:i+31} < VSR[XA]{i} || VSR[XB]{i+1:i+31}
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

For each vector element i from 0 to 3, do the following. The contents of bit 0 of word element \(i\) of VSR[XA] are concatenated with the contents of bits 1:31 of word element i of \(\mathrm{VSR}[\mathrm{XB}]\) and placed into word element i of \(\mathrm{VSR}[\mathrm{XT}]\).

Special Registers Altered
None
\begin{tabular}{ll}
\hline Extended Mnemonic & Equivalent To \\
\hline xvmovsp \(\quad \mathrm{XT}, \mathrm{XB}\) & xvcpsgnsp \(\mathrm{XT}, \mathrm{XB}, \mathrm{XB}\) \\
\hline Table 109: & \\
\hline
\end{tabular}

\section*{VSR Data Layout for xvcpsgnsp}
src1 = VSR[XA]
\begin{tabular}{|c|c|c|c|}
\hline\(S P\) & \(S P\) & \(S P\) & \(S P\) \\
\hline
\end{tabular}
src2 = VSR[XB]
\begin{tabular}{|c|c|c|c|}
\hline\(S P\) & \(S P\) & \(S P\) & \(S P\) \\
\hline
\end{tabular}
\(\operatorname{tgt}=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|l|l|l|}
\hline SP & SP & & SP & \multicolumn{2}{|c|}{ SP } \\
\hline 0 & 32 & 64 & 96 & 127 \\
\hline
\end{tabular}

\section*{Version 3.0}

VSX Vector round Double-Precision to
single-precision and Convert to
Single-Precision format XX2-form
xvcvdpsp XT,XB

\begin{tabular}{ll} 
XT & \(\leftarrow T X \| T\) \\
XB & \(\leftarrow B X \| B\) \\
ex_flag & \(\leftarrow 0 b 0\)
\end{tabular}
do \(\mathrm{i}=0\) to 127 by 64 reset_xflags()
src \(\leftarrow\) VSR[XB]\{i:i+63\}
result \(\{\mathrm{i}: \mathrm{i}+31\} \leqslant\) RoundToSP(RN, src)
result\{i+32:i+63\} \(\leftarrow\) 0xUUUU_UUUU
if(vxsnan_flag) then SetFX(VXSNAN)
if(ox_flag) then SetFX(0X)
if(ux_flag) then SetEX(UX)
if(xx_flag) then SetFX(XX)
ex_flag \(\leftarrow\) ex_flag | (VE \& vxsnan_flag)
ex_flag \(\leftarrow\) ex_flag | ( \(0 \mathrm{E} \&\) ox_flag)
ex_flag \(\leftarrow\) ex_flag | (UE \& ux_flag)
ex_flag \(\leftarrow\) ex_flag | (XE \& xx_flag)
end
if( ex_flag = 0 ) then VSR \([X T] \leftarrow\) result
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element i from 0 to 1 , do the following.
Let src be the double-precision floating-point operand in doubleword element \(i\) of VSR[XB].
src is rounded to single-precision using the rounding mode specified by RN.

The result is placed into bits 0:31 of doubleword element i of \(\operatorname{VSR}[\mathrm{XT}]\) in single-precision format.

The contents of bits 32:63 of doubleword element \(i\) of VSR[XT] are undefined.

If a trap-enabled exception occurs in any element of the vector, no results are written to \(\operatorname{VSR}[X T]\).

\section*{Special Registers Altered}

FX OX UX XX VXSNAN

VSR Data Layout for xvcvdpsp
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & DP \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|l|l|}
\hline SP & undefined & & SP \\
\hline
\end{tabular}

\section*{VSX Vector truncate Double-Precision to integer and Convert to Signed Integer Doubleword format with Saturate XX2-form}

\section*{xvcvdpsxds XT,XB}


if ( ex_flag = 0 ) then \(\operatorname{VSR}[\mathrm{XT}] \leftarrow\) result
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element \(i\) from 0 to 1 , do the following. Let \(\operatorname{src}\) be the double-precision floating-point operand in doubleword element \(i\) of VSR[ XB].

If \(\operatorname{src}\) is a NaN, the result is the value \(0 \times 8000 \_0000 \_0000 \_0000\) and VXCVI is set to 1. If src is an SNaN, VXSNAN is also set to 1 .

Otherwise, src is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than \(2^{63} \cdot 1\), the result is \(\mathrm{OXPFFF}_{\_}\)FFFF_FFF_FFFF and VXCVI is set to 1.

Otherwise, if the rounded value is less than \(-2^{63}\), the result is \(0 \times 8000,0000 \_0000 \_0000\) and VXCVI is set to 1.

Otherwise, the result is the rounded value converted to 64-bit signed-integer format, and if the result is inexact (i.e., not equal to \(\operatorname{src}\) ), \(X X\) is set to 1.

The result is placed into doubleword element i of VSR[XT].

See Table 110.
If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

\section*{Special Registers Altered}

FX XX VXSNAN VXCVI

VSR Data Layout for xvcvdpsxds
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & DP \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|}
\hline & SD \\
\hline 0 & 64 \\
\hline
\end{tabular}

\section*{Programming Note}
xvcvdpsxds rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Double-Precision Integer instruction that corresponds to the desired rounding mode, including xvrdpic which uses the rounding mode specified by the RN.

\section*{Version 3.0}


Table 110.Actions for xvcvdpsxds

\section*{VSX Vector truncate Double-Precision to integer and Convert to Signed Integer Word format with Saturate XX2-form}
xvcvdpsxws
XT, XB



Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element \(i\) from 0 to 1 , do the following. Let src be the double-precision floating-point operand in doubleword element \(i\) of VSR[ XB] .

If src is a NaN , the result is the value \(0 \times 8000 \_0000\) and VXCVI is set to 1 . If src is an SNaN, VXSNAN is also set to 1 .

Otherwise, src is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than \(2^{31} \cdot 1\), the result is \(0 \times 7 \mathrm{FFF}\) FFFF and VXCVI is set to 1 .

Otherwise, if the rounded value is less than \(\cdot 2^{31}\), the result is \(0 \times 8000.0000\) and VXCVI is set to 1 .

Otherwise, the result is the rounded value converted to 32-bit signed-integer format, and if the result is inexact (i.e., not equal to \(\operatorname{src}\) ), \(X X\) is set to 1.

The result is placed into bits 0:31 of doubleword element i of VSR[ XT].

The contents of bits 32:63 of doubleword element 1 of VSR[ XT] are undefined.

See Table 111.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT].

\section*{Special Registers Altered \\ FX XX VXSNAN VXCVI}

VSR Data Layout for xvcvdpsxws
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & DP \\
\hline
\end{tabular}
\(\operatorname{tgt}=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|l|l|}
\hline SW & undefined & SW & \multicolumn{2}{|c|}{ undefined } \\
\hline 0 & 32 & 64 & 96
\end{tabular}

\section*{Programming Note}
xvcevdpsxws rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Double-Precision Integer instruction that corresponds to the desired rounding mode, including xvrdpic which uses the rounding mode specified by RN.

\section*{Version 3.0}


Table 111.Actions for xvcvdpsxws

\section*{VSX Vector truncate Double-Precision to integer and Convert to Unsigned Integer Doubleword format with Saturate XX2-form}
\[
\text { xvcvdpuxds } \quad \mathrm{XT}, \mathrm{XB}
\]
\begin{tabular}{|l|ll|l|l|l|l|l|}
\hline 60 & & T & & III & & B & \\
\hline
\end{tabular}
\begin{tabular}{ll}
\(X T\) & \(\leftarrow T X \| T\) \\
\(X B\) & \(\leftarrow B X \| B\) \\
ex_flag & \(\leftarrow 0 b 0\)
\end{tabular}
do \(\mathrm{i}=0\) to 127 by 64
reset_xflags()
result \(\{i: 1+63\} \leftarrow\) ConvertDPtoUD (VSR [XB] \(\{i: i+63\})\)
if (vxsnan_flag) then SetFX(VXSNAN)
if (vxcvi_flag) then SetFX(vxcvi)
if (xx_flag) then SetFX(xX)
ex_flag \(\quad \leftarrow\) ex_flag | (VE \& vxsnan_flag)
(VE \& vxcvi_flag)
(XE \& xx_flag)
end
if ( ex_flag = 0 ) then VSR[XT] \(\leftarrow\) result
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element \(i\) from 0 to 1 , do the following. Let \(\operatorname{src}\) be the double-precision floating-point operand in doubleword element \(i\) of \(\operatorname{VSR}[\mathrm{XB}]\).

If \(\operatorname{src}\) is a NaN, the result is the value \(0 \times 0000 \_0000 \_0000 \_0000\) and VXCVI is set to 1. If src is \({ }^{-}\)an \(\mathrm{SNaN}, \overline{\mathrm{V}} \mathrm{XSNAN}\) is also set to 1 .

Otherwise, src is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than \(2^{64} \cdot 1\), the result is \(\mathrm{OXFFFF}_{\_}\)FFFF_FFF_FFFF and VXCVI is set to 1.

Otherwise, if the rounded value is less than 0 , the result is \(0 \times 0000,0000 \_0000 \_0000\) and VXCVI is set to 1.

Otherwise, the result is the rounded value converted to 64-bit unsigned-integer format, and if the result is inexact (i.e., not equal to \(\operatorname{src}\) ), \(X X\) is set to 1.

The result is placed into doubleword element \(i\) of VSR[ XT].

See Table 112.
If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

\section*{Special Registers Altered}

FX XX VXSNAN VXCVI

VSR Data Layout for xvcvdpuxds
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|r|l|}
\hline DP & DP \\
\hline tgt = VSR[XT] & \\
\hline \begin{tabular}{|l|l|}
\hline UD & 64 \\
\hline 0
\end{tabular}
\end{tabular}

\section*{Programming Note}
xvcvdpuxds rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Double-Precision Integer instruction that corresponds to the desired rounding mode, including xvrdpic which uses the rounding mode specified by the RN.

\section*{Version 3.0}


Table 112.Actions for xvcvdpuxds

\section*{VSX Vector truncate Double-Precision to integer and Convert to Unsigned Integer Word format with Saturate XX2-form}
\[
\text { xvcvdpuxws } \quad \mathrm{XT}, \mathrm{XB}
\]



Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element \(i\) from 0 to 1 , do the following. Let src be the double-precision floating-point operand in doubleword element \(i\) of VSR[ XB].

If src is a NaN , the result is the value \(0 \times 8000-0000\) and VXCVI is set to 1 . If src is an SNaN, VXSNAN is also set to 1.

Otherwise, src is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than \(2^{32} \cdot 1\), the result is \(0 \times F F F F \_F F F F\) and \(V X C V I\) is set to 1 .

Otherwise, if the rounded value is less than 0 , the result is \(0 \times 0000 \_0000\) and VXCVI is set to 1 .

Otherwise, the result is the rounded value converted to 32-bit unsigned-integer format, and if the result is inexact (i.e., not equal to src ), XX is set to 1.

The result is placed into bits 0:31 of doubleword element i of VSR[ XT].

The contents of bits 32:63 of doubleword element i of VSR[ XT] are undefined.

See Table 113.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

\section*{Special Registers Altered \\ FX XX VXSNAN VXCVI}

VSR Data Layout for xvcvdpuxws
\(\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]\)
\begin{tabular}{|c|c|}
\hline DP & DP \\
\hline
\end{tabular}
\(\operatorname{tgt}=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|l|l|l|l|}
\hline UW & undefined & UW & undefined \\
\hline 0 & 32 & 64 & 96
\end{tabular}

\section*{Programming Note}
xvcvdpuxws rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Double-Precision Integer instruction that corresponds to the desired rounding mode, including xvrdpic which uses the rounding mode specified by RN.

\section*{Version 3.0}


Table 113.Actions for xvcvdpuxws

VSX Vector Convert Half-Precision format to Single-Precision format XX2-form
```

xvcvhpsp XT,XB

```

```

if MSR.VSX=0 then VSX_Unavailable(l)
reset_flags()
do i =0 to 3
sIC \leftarrowbfp_CONVERT_fROM_BFP16(VSR[BX\times32 +B], word[i], hword[1])
if src.class.SNaN=1 then
result.word[i] \& bfp_CONVERT_TO_BFP32(bfp_QUIET(src))
else
result.word[i]}\leftarrow\mathrm{ bfp_CONVERT_TO_BFP32(src)
vxsnan_flag \leftarrow src.class.SNaN
if(vxsnan_flag) then SetFX(FPSCR. VXSNAN)
ex_flag\leftarrowex_flag|(FPSCR.VE \& vxsnan_flag)
end
if ex_flag=0 then VSR[XT] \leftarrowresult

```

Let XT be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each integer value i from 0 to 3 , do the following. Let \(\operatorname{src}\) be the half-precision floating-point value in the rightmost halfword of word element i of VSR[ XB].

If \(s r^{c}\) is an SNaN , the result is the single-precision representation of that SNaN converted to a QNaN .

Otherwise, if src is a QNaN, the result is the single-precision representation of that QNaN.

Otherwise, if src is an Infinity, the result is the single-precision representation of Infinity with the same sign as sic.

Otherwise, if src is a Zero, the result is the single-precision representation of Zero with the same sign as sic.

Otherwise, if src is a denormal value, the result is the normalized single-precision representation of src.

Otherwise, the result is the single-precision representation of src .

The result is placed into word element i of VSR[ XT].

If a trap-enabled exception occurs, VSR[XT] is not modified.

Special Registers Altered:
FX VXSNAN

VSR Data Layout for xvcvhpsp


\section*{Version 3.0}

VSX Vector Convert Single-Precision to Double-Precision format XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline xvcvspd & & XT,XB & & & & \\
\hline \[
{ }_{0} 60
\] & \[
{ }_{6} \quad \mathrm{~T}
\] & \[
\left.\right|_{11} I I I
\] & \(1{ }_{16}\) B & 21 & & \(\left|\begin{array}{c}\text { BX TXX } \\ 30 \\ 31\end{array}\right|\) \\
\hline \multicolumn{7}{|l|}{\(X T \quad \leftarrow T X \| T\)} \\
\hline \multicolumn{7}{|l|}{} \\
\hline \multicolumn{7}{|l|}{ex_flag \(\leftarrow\) Obo} \\
\hline \multicolumn{7}{|l|}{do \(i=0\) to 127 by 64} \\
\hline \multicolumn{7}{|c|}{reset_xflags ()} \\
\hline \multicolumn{7}{|c|}{result \(\{i: i+63\} \leftarrow\) ConvertSPtoDP (VSR \([\mathrm{XB}]\{i: i+31\})\)} \\
\hline \multicolumn{7}{|c|}{if (vxsnan_flag) then SetFX (VXSNAN)} \\
\hline \multicolumn{7}{|c|}{ex_flag \(\leftarrow\) ex_flag | (VE \& vxsnan_flag)} \\
\hline \multicolumn{7}{|l|}{end} \\
\hline \multicolumn{7}{|l|}{if ( ex_flag = 0 ) then VSR [XT] \(\leftarrow\) result} \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each vector element \(i\) from 0 to 1 , do the following.
Let src be the single-precision floating-point operand in bits 0:31 of doubleword element i of VSR[ XB].

SrC is placed into doubleword element i of V SR[XT] in double-precison format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

Special Registers Altered
FX VXSNAN

VSR Data Layout for xvcvspdp
src \(=\) VSR[XB]
\begin{tabular}{|c|c|c|c|}
\hline SP & unused & SP & unused \\
\hline
\end{tabular}
tgt \(=\mathrm{VSR}[\mathrm{XT}]\)
\begin{tabular}{|c|ccc|}
\hline DP & DP \\
\hline 0 & 32 & 64 & 96 \\
\hline
\end{tabular}

\section*{VSX Vector round and Convert Single-Precision format to Half-Precision format XX2-form}
```

xvcvsphp XT,XB

```

```

if MSR.VSX=O then VSX_Unavailable(l

```
if MSR.VSX=O then VSX_Unavailable(l
reset_flags()
reset_flags()
do i =0 to 3
do i =0 to 3
    sic <bfp_CONVERT_FROM_BFP32[VSR[BX\times32+B],word[i])
    sic <bfp_CONVERT_FROM_BFP32[VSR[BX\times32+B],word[i])
    rnd }\leftarrow\mathrm{ bfp_ROUND_TO_BFP16(FPSCR,RN,rnd)
    rnd }\leftarrow\mathrm{ bfp_ROUND_TO_BFP16(FPSCR,RN,rnd)
    result.hword[2xi]}\leftarrow0\times000
    result.hword[2xi]}\leftarrow0\times000
    result.hword[2xi +1] &bfp_CONVERT_TO_BFP16(rnd)
    result.hword[2xi +1] &bfp_CONVERT_TO_BFP16(rnd)
    if(vxsnan_flag) then SetFX(FPSCR. VXSNAN)
    if(vxsnan_flag) then SetFX(FPSCR. VXSNAN)
    if(Ox_flag) then SetFX(FPSCR.OX)
    if(Ox_flag) then SetFX(FPSCR.OX)
    if(ux_flag) then SetFX(FPSCR.UX)
    if(ux_flag) then SetFX(FPSCR.UX)
    if(xx_flag) then SetFX(FPSCR. XX)
    if(xx_flag) then SetFX(FPSCR. XX)
    ex_flag \leftarrowex_flag|(FPSCR.VE & vxsnan_flag)
    ex_flag \leftarrowex_flag|(FPSCR.VE & vxsnan_flag)
        | (FPSCR.OE & Ox_flag)
        | (FPSCR.OE & Ox_flag)
        | (FPSCR.UE & ux_flag)
        | (FPSCR.UE & ux_flag)
        | (FPSCR.XE & XX_flag)
        | (FPSCR.XE & XX_flag)
end
```

if(ex_flag=0) then VSR $[X T] \leftarrow$ result

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each integer value i from 0 to 3 , do the following. Let src be the half-precision floating-point value represented in single-precision format in word element i of VSR[ XB].

If $s r^{\circ}$ is an SNaN , the result is the half-precision representation of that SNaN converted to a QNaN .

Otherwise, if src is a QNaN , the result is the half-precision representation of that QNaN .

Otherwise, if src is an Infinity, the result is the half-precision representation of Infinity with the same sign as sic.

Otherwise, if $\operatorname{src}$ is a Zero, the result is the half-precision representation of Zero with the same sign as src.

Otherwise, the result is the half-precision representation of $\operatorname{src}$ rounded to half-precision using the rounding mode specified by RN.

The result is zero-extended and placed into word element i of VSR[XT].

If a trap-enabled exception occurs, VSR[XT] is not modified.

Special Registers Altered:
FX VXSNAN OX UX XX

## VSR Data Layout for xvcvsphp



## Version 3.0

## VSX Vector truncate Single-Precision to integer and Convert to Signed Integer Doubleword format with Saturate XX2-form

xvcvspsxds XT,XB

| 60 |  | T |  | III |  | B |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  |  | 408 |  |


| XT | $\leftarrow \mathbb{T X} \\| \mathbb{T}$ |
| :--- | :--- |
| XB | $\leftarrow \mathrm{BX} \\| \mathrm{B}$ |
| ex_flag | $\leftarrow 0 \mathrm{bO}$ |

do $i=0$ to 127 by 64
reset_xflags()
result $\{i: i+63\} \leftarrow$ ConvertSPtoSD (VSR $[\mathrm{XB}]\{i: i+31\})$
if (vxsnan_flag) then SetFX (VXSNAN)
if (vxcvi_flag) then SetFX(VXCVI)
if (xx_flag) then $\operatorname{SetFX}(X X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
| (VE \& vxcvi_flag)
| (XE \& xx_flag)
end
if ( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element ifrom 0 to 1, do the following. Let SrC be the single-precision floating-point operand in word element $i \times 2$ of VSR[ XB].

If $s r c$ is a $N a N$, the result is the value $0 \times 8000-0000-0000-0000$ and VXCVI is set to 1. If src is an SNaN, VXSNAN is also set to 1 .

Otherwise, src is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than $2^{63} \cdot 1$, the result is $0 \times 7 F F F$ _FFF_FFFF_FFFF and VXCVI is set to 1.

Otherwise, if the rounded value is less than $\cdot 2^{63}$, the result is $0 \times 8000 \_0000 \_0000 \_0000$ and VXCVI is set to 1.

Otherwise, the result is the rounded value converted to 64-bit signed-integer format, and if the result is inexact (i.e., not equal to $\operatorname{src}$ ), $X X$ is set to 1.

The result is placed into doubleword element $i$ of VSR[ XT].

See Table 113.
If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered

FX XX VXSNAN VXCVI

## VSR Data Layout for xvcvspsxds

$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| $S P$ | unused | SP | unused |
| :---: | :---: | :---: | :---: |

tgt $=\operatorname{VSR}[\mathrm{XT}]$

|  | SD | SD |
| :---: | :---: | :---: |
| 0 | 32 | 64 |

## Programming Note

xvcuspsxds rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Single-Precision Integer instruction that corresponds to the desired rounding mode, including xurspic which uses the rounding mode specified by RN.

|  | 山 | 山 |  | Returned Results and Status Setting |
| :---: | :---: | :---: | :---: | :---: |
| SrC $\leq$ Nmin-1 | 0 | - | - | T(Nmin), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| Nmin-1 < SrC < Nmin | - | 0 | yes | $\mathrm{T}(\mathrm{Nmin}), \mathrm{fx}(\mathrm{XX})$ |
|  |  | 1 | yes | fx(XX), error() |
| src $=$ Nmin | - | - | no | T(Nmin) |
| Nmin < SrC < Nmax | - | - | no | T(ConvertSPtoSD(RoundToSPintegerTrunc(src))) |
|  |  | 0 | yes | T(ConvertSPtoSD(RoundToSPintegerTrunc(src))), fx(XX) |
|  |  | 1 | yes | fx(XX), error() |
| SrC $=$ Nmax | - | - | no | T(Nmax) <br> Note: This case cannot occur as Nmax is not representable in SP format but is included here for completeness. |
| Nmax < src < Nmax+1 | - | 0 | yes | T(Nmax), fx(XX) |
|  |  | 1 | yes | fx(XX), error() |
| Src $\geq$ Nmax+1 | 0 | - | - | T(Nmax), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| src is a QNaN | 0 | - | - | T(Nmin), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| src is a SNaN | 0 | - | - | T(Nmin), fx(VXCVI), fx(VXSNAN) |
|  | 1 | - | - | fx(VXCVI), fx(VXSNAN), error() |
| Explanation: |  |  |  |  |
| fx(x) | $F X$ is set to 1 if $\mathrm{x}=0 . \mathrm{X}$ is set to 1 . |  |  |  |
| error() | The system error handler is invoked for the trap-enabled exception if the FEO and FE1 bits in the Machine State Register are set to any mode other than the ignore-exception mode. |  |  |  |
|  | Update of VSR[XT] is suppressed. |  |  |  |
| N mi $n$ | The smallest signed integer doubleword value, - $2^{63}\left(0 \times 8000 \_0000 \_0000 \_0000\right)$. |  |  |  |
| $N$ max | The largest signed integer doubleword value, $2^{63} .1$ ( $0 \times 7$ FFF_FFFF_ FFFF_ FFFF). |  |  |  |
| src | The single-precision floating-point value in word element i of VSR $[X B]$ (where i $\in\{0,2\}$ ). |  |  |  |
| T ( x ) | The signed integer doubleword value $x$ is placed in doubleword element $i$ of VSR[ $X T]$ (where $i \in\{0,1\}$ ). |  |  |  |

Table 114.Actions for xvcvspsxds

## Version 3.0

VSX Vector truncate Single-Precision to integer and Convert to Signed Integer Word format with Saturate XX2-form

$$
\text { xvcvspsxws } \quad \text { XT,XB }
$$

| 60 |  | T |  | III |  | B |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  | 15 |


| XT | $\leftarrow \mathbb{T X} \\| \mathbb{T}$ |
| :--- | :--- |
| XB | $\leftarrow \mathrm{BX} \\| \mathrm{B}$ |
| ex_flag | $\leftarrow 0 \mathrm{bO}$ |

do $i=0$ to 127 by 32
reset_xflags ()
result $\{i: 1+31\} \leftarrow$ ConvertSPToSW(VSR [XB] $\{i: i+31\})$
if (vxsnan_flag) then SetFX (vxsnaiv)
if (vxcvi_flag) then SetFX(vxcvi)
if (xx_flag) then Setex (xX)
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
| (VE \& vxcvi_flag)
| (XE \& xx_flag)
end
if ( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element ifrom 0 to 3 , do the following. Let SrC be the single-precision floating-point operand in word element $i$ of VSR[ XB].

If $\operatorname{src}$ is a NaN , the result is the value $0 \times 8000.0000$ and VXCVI is set to 1 . If src is an SNaN, VXSNAN is also set to 1 .

Otherwise, src is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than $2^{31} \cdot 1$, the result is $0 \times 7 F F F \_F F F$, and $V X C V I$ is set to 1 .

Otherwise, if the rounded value is less than $\cdot 2^{31}$, the result is $0 \times 8000 \_0000$, and VXCVI is set to 1 .

Otherwise, the result is the rounded value converted to 32-bit signed-integer format, and if the result is inexact (i.e., not equal to $s i c$ ), $X X$ is set to 1.

The result is placed into word element i of VSR[XT].

See Table 113.
If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered

FX XX VXSNAN VXCVI

## VSR Data Layout for xvcvspsxws

$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

| SW | SW | SW | SW |  |
| :--- | :--- | :--- | :--- | :---: |
| 0 | 32 | 64 | 96 |  |

## Programming Note

xvcvspsxws rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Single-Precision Integer instruction that corresponds to the desired rounding mode, including xvrspic which uses the rounding mode specified by RN.

|  | ш | 山 |  | Returned Results and Status Setting |
| :---: | :---: | :---: | :---: | :---: |
| Src $\leq$ Nmin-1 | 0 | - | - | T(Nmin), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| Nmin-1 < SrC < Nmin | - | 0 | yes | T(Nmin), fx(XX) |
|  |  | 1 | yes | fx(XX), error() |
| Src $=$ Nmin | - | - | no | T(Nmin) |
| Nmin < Src < Nmax | - | - | no | T(ConvertSPtoSW(RoundToSPintegerTrunc(src))) |
|  |  | 0 | yes | T(ConvertSPtoSW(RoundToSPintegerTrunc(src))), fx(XX) |
|  |  | 1 | yes | fx(XX), error() |
| SrC $=$ Nmax | - | - | no | T(Nmax) <br> Note: This case cannot occur as Nmax is not representable in SP format but is included here for completeness. |
| Nmax < SrC < Nmax+1 | - | 0 | yes | T(Nmax), fx(XX) |
|  |  | 1 | yes | $f x(X X)$, error() |
| $\operatorname{SrC} \geq \mathrm{Nmax}+1$ | 0 | - | - | T(Nmax), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| src is a QNaN | 0 | - | - | T(Nmin), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| src is a SNaN | 0 | - | - | T(Nmin), fx(VXCVI), fx(VXSNAN) |
|  | 1 | - | - | fx(VXCVI), fx(VXSNAN), error() |
| Explanation: |  |  |  |  |
| fx(x) | $F X$ is set to 1 if $\mathrm{x}=0 . \mathrm{x}$ is set to 1 . |  |  |  |
| error() | The system error handler is invoked for the trap-enabled exception if the FEO and FE1 bits in the Machine State Register are set to any mode other than the ignore-exception mode. |  |  |  |
|  | Update of VSR[XT] is suppressed. |  |  |  |
| Nmin | The smallest signed integer word value, $-2^{31}$ (0x8000_0000). |  |  |  |
| Nmax | The largest signed integer word value, $2^{31}-1$ (0x7FFF_FFFF). |  |  |  |
| src | The single-precision floating-point value in word element i of VSR[XB] (where $i \in\{0,1,2,3\}$ ). |  |  |  |
| T(x) | The signed integer word value $x$ is placed in word element $i$ of VSR[XT] (where $i \in\{0,1,2,3\}$ ). |  |  |  |

Table 115.Actions for xvcvspsxws

## Version 3.0

## VSX Vector truncate Single-Precision to integer and Convert to Unsigned Integer Doubleword format with Saturate XX2-form

xvcvspuxds $\quad X T, X B$


| XT | $\leftarrow \mathbb{T X} \\| \mathbb{T}$ |
| :--- | :--- |
| XB | $\leftarrow \mathrm{BX} \\| \mathrm{B}$ |
| ex_flag | $\leftarrow 0 \mathrm{bO}$ |

do $\mathrm{i}=0$ to 127 by 64
reset_xflags ()
result $\{\mathrm{i}: \mathrm{i}+63\} \leftarrow$ ConvertSPtouD (VSR [XB] $\{i: i+31\})$
if (vxsnan_flag) then SetFX (vxsnain)
if (vxcvi_flag) then SetFX(vxcvi)
if (xx_flag) then $\operatorname{SetFX}(x X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
| (VE \& vxcvi_flag)
| (XE \& xx_flag)
end
if ( ex_flag =0) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let SrC be the single-precision floating-point operand in word element $i \times 2$ of VSR[ XB].

If src is a NaN, the result is the value $0 \times 0000-0000-0000-0000$ and VXCVI is set to 1. If src is an SNaN, VXSNAN is also set to 1 .

Otherwise, src is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than $2^{64} \cdot 1$, the result is $\mathrm{OXFFFF}^{2}$ FFFF_FFFF_FFFF and VXCVI is set to 1.

Otherwise, if the rounded value is less than 0 , the result is $0 \times 0000 \_0000 \_0000 \_0000$ and VXCVI is set to 1.

Otherwise, the result is the rounded value converted to 64-bit unsigned-integer format, and if the result is inexact (i.e., not equal to $\operatorname{src}$ ), $X X$ is set to 1.

The result is placed into doubleword element $i$ of VSR[XT].

See Table 113.
If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered

FX XX VXSNAN VXCVI

## VSR Data Layout for xvcvspuxds

$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| $S P$ | unused | SP | unused |
| :---: | :---: | :---: | :---: |

tgt $=\operatorname{VSR}[\mathrm{XT}]$

|  | UD | UD |  |
| :---: | :---: | :---: | :---: |
| 0 | 32 | 64 | 96 |

## Programming Note

xvcvspuxds rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Single-Precision Integer instruction that corresponds to the desired rounding mode, including xurspic which uses the rounding mode specified by RN.

|  | ш | 岗 |  | Returned Results and Status Setting |
| :---: | :---: | :---: | :---: | :---: |
| Src $\leq$ Nmin-1 | 0 | - | - | T(Nmin), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| Nmin-1 < SrC < Nmin | - | 0 | yes | T(Nmin), fx(XX) |
|  |  | 1 | yes | fx(XX), error() |
| Src $=$ Nmin | - | - | no | T(Nmin) |
| Nmin < SrC < Nmax | - | - | no | T(ConvertSPtoUD(RoundToSPintegerTrunc(src))) |
|  |  | 0 | yes | T(ConvertSPtoUD(RoundToSPintegerTrunc(src))), fx(XX) |
|  |  | 1 | yes | fx(XX), error() |
| SrC $=$ Nmax | - | - | no | T(Nmax) <br> Note: This case cannot occur as Nmax is not representable in SP format but is included here for completeness. |
| Nmax < src < Nmax+1 | - | 0 | yes | T(Nmax), fx(XX) |
|  |  | 1 | yes | $f x(X X)$, error() |
| Src $\geq \mathrm{Nmax}+1$ | 0 | - | - | T(Nmax), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| src is a QNaN | 0 | - | - | T(Nmin), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| src is a SNaN | 0 | - | - | T(Nmin), fx(VXCVI), fx(VXSNAN) |
|  | 1 | - | - | fx(VXCVI), fx(VXSNAN), error() |
| Explanation: |  |  |  |  |
| fx(x) | FX is set to 1 if $\mathrm{x}=0 . \mathrm{x}$ is set to 1 . |  |  |  |
| error() | The system error handler is invoked for the trap-enabled exception if the FEO and FE1 bits in the Machine State Register are set to any mode other than the ignore-exception mode. |  |  |  |
| Nmin | The smallest unsigned integer doubleword value, 0 (0x00000_0000_0000_0000). |  |  |  |
| Nmax | The largest unsigned integer doubleword value, $2^{64}-1$ (0xFFFF_FFFF_FFFF_FFFF). |  |  |  |
| src | The single-precision floating-point value in word element $i$ of VSR[XB] (where $i \in\{0,2\}$ ). |  |  |  |
| T(x) | The unsigned integer doubleword value $x$ is placed in doubleword element $i$ of VSR[XT] (where $i \in\{0,1\}$ ). |  |  |  |

Table 116.Actions for xvcvspuxds

## Version 3.0

VSX Vector truncate Single-Precision to integer and Convert to Unsigned Integer Word format with Saturate XX2-form
xvcvspuxws XT,XB


| XT | $\leftarrow$ TX \\| |
| :--- | :--- |
| XB | $\leftarrow$ BX \\|B |
| ex_flag | $\leftarrow 0 b 0$ |

do i=0 to 127 by 32
reset_xflags()
result $\{i: i+31\} \leftarrow$ ConvertSPtoUW(VSR $[X B]\{i: i+31\})$
if(vxsnan_flag) then SetFX(vXSNAN)
if(vxcvi_flag) then SetEX(vXCVI)
if(xx_flag) then $\operatorname{SetFX}(X X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
| (VE \& vxcvi_flag)
| (XE \& xx_flag)
end
if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element ifrom 0 to 3 , do the following. Let src be the single-precision floating-point operand in word element $i$ of VSR[ XB].

If srC is a NaN , the result is the value $0 \times 0000_{\mathbf{Z}} 0000$ and VXCVI is set to 1 . If src is an SNaN, VXSNAN is also set to 1 .

Otherwise, SrC is rounded to a floating-point integer using the rounding mode Round Toward Zero.

If the rounded value is greater than $2^{32} \cdot 1$, the result is $\mathrm{OXFFFF}_{-} \mathrm{FFFF}$ and VXCVI is set to 1 .

Otherwise, if the rounded value is less than 0 , the result is $0 \times 0000 \_0000$ and VXCVI is set to 1 .

Otherwise, the result is the rounded value converted to 32-bit unsigned-integer format, and if the result is inexact (i.e., not equal to $s r^{\prime}$ ), $X X$ is set to 1 .

The result is placed into word element i of VSR[XT].

See Table 113.
If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

Special Registers Altered
FX XX VXSNAN VXCVI

## VSR Data Layout for xvcvspuxws

$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

| UW | UW | UW | UW |  |
| :--- | :--- | :--- | :--- | :---: |
| 0 | 32 | 64 | 96 |  |

## Programming Note

xvcvspuxws rounds using Round towards Zero rounding mode. For other rounding modes, software must use a Round to Single-Precision Integer instruction that corresponds to the desired rounding mode, including xvrspic which uses the rounding mode specified by RN.

|  | ш | 岗 | Inexact? ( RoundToSPintegerTrunc(src) $\neq \mathrm{src}$ ) | Returned Results and Status Setting |
| :---: | :---: | :---: | :---: | :---: |
| Src $\leq$ Nmin-1 | 0 | - | - | T(Nmin), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| Nmin-1 < SrC < Nmin | - | 0 | yes | T(Nmin), fx(XX) |
|  |  | 1 | yes | fx(XX), error() |
| Src $=$ Nmin | - | - | no | T(Nmin) |
| Nmin < SrC < Nmax | - | - | no | T(ConvertSPtoUW(RoundToSPintegerTrunc(src))) |
|  |  | 0 | yes | T(ConvertSPtoUW(RoundToSPintegerTrunc(src))), fx(XX) |
|  |  | 1 | yes | fx(XX), error() |
| SrC $=$ Nmax | - | - | no | T(Nmax) <br> Note: This case cannot occur as Nmax is not representable in SP format but is included here for completeness. |
| Nmax < src < Nmax+1 | - | 0 | yes | T(Nmax), fx(XX) |
|  |  | 1 | yes | $f x(X X)$, error() |
| Src $\geq \mathrm{Nmax}+1$ | 0 | - | - | T(Nmax), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| src is a QNaN | 0 | - | - | T(Nmin), fx(VXCVI) |
|  | 1 | - | - | fx(VXCVI), error() |
| src is a SNaN | 0 | - | - | T(Nmin), fx(VXCVI), fx(VXSNAN) |
|  | 1 | - | - | fx(VXCVI), fx(VXSNAN), error() |
| Explanation: |  |  |  |  |
| fx(x) | $F X$ is set to 1 if $\mathrm{x}=0 . \mathrm{x}$ is set to 1 . |  |  |  |
| error() | The system error handler is invoked for the trap-enabled exception if the FEO and FE1 bits in the Machine State Register are set to any mode other than the ignore-exception mode. |  |  |  |
| Nmin | The smallest unsigned integer word value, 0 (0x0000_0000). |  |  |  |
| Nmax | The largest unsigned integer word value, $2^{32}-1$ (0xFFFF_FFFF). |  |  |  |
| src | The single-precision floating-point value in word element $i$ of VSR[XB] (where $i \in\{0,1,2,3\}$ ). |  |  |  |
| T(x) | The unsigned integer word value $x$ is placed in word element $i$ of VSR[XT] (where $i \in\{0,1,2,3\}$ ). |  |  |  |

Table 117.Actions for xvcvspuxws

VSX Vector Convert and round Signed Integer Doubleword to Double-Precision format XX2-form

| xvcvsxddp XT,XB |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll}  & 60 \\ 0 \end{array}$ | ${ } \quad \mathrm{T}$ | ${ }_{11} \quad \text { III }$ | ${ }_{16} \text { B }$ | 21 | 504 | BX <br> 30 <br> 30 <br> 31 |
| XT $\leftarrow T X \\| T$ |  |  |  |  |  |  |
| $X B \leqslant B X \\| B$ |  |  |  |  |  |  |
| ex_flag $\leftarrow 0$ 00 |  |  |  |  |  |  |
| do i=0 to 127 by 64 |  |  |  |  |  |  |
| reset_xflags() |  |  |  |  |  |  |
| v 0:inf $\} \quad \leftarrow$ ConvertSDtoFP(VSR[XB]\{i:i+63 |  |  |  |  |  |  |
| result $\{\mathrm{i}: 1+63\} \leftarrow$ RoundToDP(RN, v) |  |  |  |  |  |  |
| if(xx_flag) then SetFX(XX) |  |  |  |  |  |  |
| ex_flag $\leftarrow$ ex_flag \| (XE \& xx_flag) |  |  |  |  |  |  |
| end |  |  |  |  |  |  |
| if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result |  |  |  |  |  |  |

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element ifrom 0 to 1 , do the following. Let src be the signed integer in doubleword element i of VSR[ XB].

Src is converted to an unbounded-precision floating-point value and rounded to double-precision using the rounding mode specified by RN.

The result is placed into doubleword element $i$ of VSR[ XT] in double-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered

FX XX


VSX Vector Convert and round Signed Integer Doubleword to Single-Precision format XX2-form
xvcvsxdsp XT,XB


| $X T \leqslant T X \\| T$ |  |
| :---: | :---: |
| $X B \quad \leftarrow B X\|\mid B$ |  |
| ex_flag $\leftarrow 0$ 00 |  |
| do i=0 to 127 by 64 |  |
| reset_xflags() |  |
| $v\{0: i n f\}$ | $\leftarrow$ ConvertSDtoFP(VSR[XB]\{i: $1+63\}$ ) |
| result $\{i: 1+31\}$ | $\leftarrow$ RoundToSP(RN, v) |
| result $\{1+32: i+63\} \leftarrow$ 0xUUUU_UUUU |  |
| if(xx_flag) then SetFX(XX) |  |
| ex_flag | $\leftarrow$ ex_flag \| (XE \& xx_flag) |
| end |  |

$$
\text { if }(\text { ex_flag }=0) \text { then } \operatorname{VSR}[X T] \leftarrow \text { result }
$$

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element $i$ from 0 to 1, do the following. Let src be the signed integer in doubleword element i of VSR[ XB].
src is converted to an unbounded-precision floating-point value and rounded to single-precision using the rounding mode specified by RN.

The result is placed into bits $0: 31$ of doubleword element $i$ of VSR[ XT] in single-precision format.

The contents of bits 32:63 of doubleword element i of VSR[ XT] are undefined.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered

FX XX
VSR Data Layout for xvcvsxdsp
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| SD | SD |
| :---: | :---: |

$\operatorname{tgt}=\mathrm{VSR}[\mathrm{XT}]$


## VSX Vector Convert Signed Integer Word to Double-Precision format XX2-form



Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let src be the signed integer in bits 0:31 of doubleword element i of VSR[XB].
src is converted to an unbounded-precision floating-point value and rounded to double-precision using the rounding mode specified by RN.

The result is placed into doubleword element $i$ of VSR[XT] in double-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

## FX XX

VSR Data Layout for xvcvsxwdp
src = VSR[XB]

| SW | unused | SW | unused |
| :---: | :---: | :---: | :---: |

$\operatorname{tgt}=\mathrm{VSR}[\mathrm{XT}]$

|  | DP | DP |
| :---: | :---: | :---: |
| 0 | 32 | 64 |

VSX Vector Convert and round Signed Integer Word to Single-Precision format XX2-form
xvcvsxwsp XT,XB

| 60 | 6 | T | 11 | III | 16 | B | 21 | 184 | BXXTX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

$$
X T \quad \leftarrow T X \| T
$$

$$
X B \quad \leftarrow B X \| B
$$

$$
\text { ex_flag } \leftarrow 0 b 0
$$

do $\mathrm{i}=0$ to 127 by 32
reset_xflags()
$v\{0:$ inf $\} \leftarrow$ ConvertSWtoFP(VSR[XB]\{i:i+31\})
result $\{\mathrm{i}: 1+31\} \leftarrow$ RoundToSP(RN, v$)$
if(xx_flag) then SetFX(XX)

$$
\text { ex_flag } \leftarrow \text { ex_flag | (XE \& xx_flag) }
$$

end
if( ex_flag $=0)$ then VSR $[X T] \leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3 , do the following. Let src be the signed integer in word element $i$ of VSR[XB].
src is converted to an unbounded-precision floating-point value and rounded to single-precision using the rounding mode specified by RN.

The result is placed into word element i of VSR[XT] in single-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to $\operatorname{VSR}[X T]$.

## Special Registers Altered

FX XX

VSR Data Layout for xvevsxwsp
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| SW | SW | SW | SW |
| :---: | :---: | :---: | :---: |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

| SP | SP | SP | SP |  |  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :---: | :---: | :---: | :---: | :---: |
| 0 | 32 | 64 |  |  |  |  |  |  | 127 |

## VSX Vector Convert and round Unsigned Integer Doubleword to Double-Precision format XX2-form

$$
\text { xvcvuxddp } \quad X T, X B
$$

| 60 | 6 | T | 11 |  | 16 | B | 21 | 488 | 3 BXT |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

```
XT}\leqslantTX|
XB }\leftarrow\textrm{BX|B
ex_flag \leftarrow 0b0
do i=0 to 127 by 64
    reset_xflags()
    v{0:inf}}\leftarrow\mathrm{ ConvertUDtoFP(VSR[XB]{i:i+63})
    result{i:i+63}}\leftarrow~RoundToDP(RN,v
    if(xx_flag) then SetFX(XX)
    ex_flag \leftarrowex_flag | (XE & xx_flag)
end
if( ex_flag = 0 ) then VSR[XT] }\leftarrow\mathrm{ result
```

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let src be the unsigned integer in doubleword element i of VSR[XB]. src is converted to an unbounded-precision floating-point value and rounded to double-precision using the rounding mode specified by RN.

The result is placed into doubleword element $i$ of $\mathrm{VSR}[\mathrm{XT}]$ in double-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to $\operatorname{VSR}[\mathrm{XT}]$.


## VSX Vector Convert and round Unsigned Integer Doubleword to Single-Precision format XX2-form

$$
\text { xvcvuxdsp } \quad \text { XT,XB }
$$



```
XT }\leftarrowTX||
XB}\leftarrow\textrm{BX|B
ex_flag \leftarrow 0b0
do i=0 to 127 by 64
    reset_xflags()
    v{0:inf}}\leftarrow~\mathrm{ ConvertUDtoFP(VSR[XB]{i:i+63})
    result{i:i+31}}\leftarrow~\mathrm{ RoundToSP(RN,v)
    result{i+32:i+63}}\leftarrow 0xUUUU_UUUU
    if(xx_flag) then SetFX(XX)
    ex_flag & ex_flag | (XE & Xx_flag)
end
```

if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let src be the unsigned integer in doubleword element i of $\operatorname{VSR}[X B]$.
src is converted to an unbounded-precision floating-point value and rounded to single-precision using the rounding mode specified by RN.

The result is placed into bits $0: 31$ of doubleword element i of VSR[XT] in single-precision format.

The contents of bits 32:63 of doubleword element $i$ of $\operatorname{VSR}[X T]$ are undefined.

If a trap-enabled exception occurs in any element of the vector, no results are written to $\operatorname{VSR}[X T]$.

## Special Registers Altered

FX XX
VSR Data Layout for xvcvuxdsp
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| UD | UD |
| :---: | :---: |

tgt $=\operatorname{VSR}[\mathrm{XT}]$


## VSX Vector Convert and round Unsigned Integer Word to Double-Precision format xx2-form

```
xvcvuxwdp XT,XB
```



Let XT be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let src be the unsigned integer in bits $0: 31$ of doubleword element i of VSR[XB].
src is converted to an unbounded-precision floating-point value and rounded to double-precision using the rounding mode specified by RN.

The result is placed into doubleword element $i$ of VSR[XT] in double-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].


## VSX Vector Convert and round Unsigned Integer Word to Single-Precision format XX2-form

```
\[
\text { xvcvuxwsp } \quad X T, X B
\]
```



Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3 , do the following. Let src be the unsigned integer in word element i of VSR[XB].
src is converted to an unbounded-precision floating-point value and rounded to single-precision using the rounding mode specified by RN.

The result is placed into word element i of $\operatorname{VSR}[X T]$ in single-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered <br> FX XX

## VSR Data Layout for xvcvuxwsp

$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| UW | UW | UW | UW |
| :---: | :---: | :---: | :---: |
| tgt = VSR[XT] |  |  |  |
| SP SP SP SP <br> 0    |  |  |  | | 32 |
| :--- |

VSX Vector Divide Double-Precision XX3-form
xvdivdp XT,XA,XB


| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

$$
\text { do i=0 to } 127 \text { by } 64
$$

reset_xflags()

$$
\operatorname{src1} \quad \leftarrow \operatorname{VSR}[X A]\{i: i+63\}
$$

$$
\operatorname{src2} \quad \leftarrow \operatorname{VSR}[X B]\{i: i+63\}
$$

$$
\text { v\{0:inf }\} \quad \leftarrow \text { DivideDP(src1,src2) }
$$

$$
\text { result }\{i: i+63\} \leftarrow \text { RoundToDP }(R N, v)
$$

if(vxsnan_flag) then SetFX(VXSNAN)
if(vxidi_flag) then SetFX(VXIDI)
if(vxisi_flag) then SetFX(VXZDZ)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)

$$
\text { if(xx_flag) then } \operatorname{SetFX}(X X)
$$

$$
\text { if(zx_flag) then } \operatorname{SetFX}(Z X)
$$

$$
\text { ex_flag } \leftarrow \text { ex_flag | (VE \& vxsnan_flag) }
$$

$$
\text { ex_flag } \quad \leftarrow \text { ex_flag | (VE \& vxidi_flag) }
$$

$$
\text { ex_flag } \quad \leftarrow \text { ex_flag | (VE \& vxzdz_flag) }
$$

$$
\text { ex_flag } \leftarrow \text { ex_flag } \mid \text { (OE \& ox_flag) }
$$

$$
\text { ex_flag } \quad \leftarrow \text { ex_flag } \mid \text { (UE \& ux_flag) }
$$

$$
\text { ex_flag } \leftarrow \text { ex_flag } \mid \text { (ZE \& zx_flag) }
$$

end

$$
\text { ex_flag } \leftarrow \text { ex_flag | (XE \& xx_flag) }
$$

if( ex_flag $=0$ ) then VSR $[X T] \leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let src1 be the double-precision floating-point operand in doubleword element $i$ of VSR[XA].

Let src2 be the double-precision floating-point operand in doubleword element $i$ of VSR[XB].
src1 is divided ${ }^{[1]}$ by src2, producing a quotient having unbounded range and precision.

The quotient is normalized ${ }^{[2]}$.
See Table 118.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

[^2]
## 700

|  | src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| -Infinity | $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{dQNaN} \\ & \text { vxidi_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{V} \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $\mathrm{V} \leftarrow-$ Infinity | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vxidi_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ |
| -NZF | $v \leftarrow+$ Zero | $\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)$ | $\mathrm{v} \leftarrow+$ Infinity Zx_flag $\leftarrow 1$ | $\begin{aligned} & \hline \text { V } \leftarrow-\text { Infinity } \\ & \text { zx_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)$ | $v \leftarrow$-Zero | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l\|} \hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| -Zero | $v \leftarrow+$ Zero | $v \leftarrow+$ Zero | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \mathrm{vxzdz} \text { _lag } \leftarrow 1 \end{aligned}$ | $\begin{array}{\|l\|} \hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\ \mathrm{vxzdz} \text { _lag } \leftarrow 1 \\ \hline \end{array}$ | $\mathrm{v} \leftarrow$-Zero | $\mathrm{v} \leftarrow$-Zero | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l\|} \hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\ \text { vxsnan_flag } \leftarrow 1 \end{array}$ |
| $\overline{\text { - }}$ +Zero | $v \leftarrow$-Zero | $v \leftarrow$-Zero | $\mathrm{v} \leftarrow \mathrm{dQNaN}$ vxzdz_flag $\leftarrow 1$ | $\mathrm{v} \leftarrow \mathrm{dQNaN}$ <br> vxzdz_flag $\leftarrow 1$ | $v \leftarrow+$ Zero | $v \leftarrow+$ Zero | $\mathrm{v} \leftarrow \mathrm{src} 2$ | $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{Q}(\text { src2 }) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ |
| あ +NZF | $\mathrm{V} \leftarrow$-Zero | $\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)$ | $\begin{aligned} & \mathrm{V} \leftarrow-\operatorname{lnfinity} \\ & \mathrm{zx} \text { _flag } \leftarrow 1 \end{aligned}$ | $\mathrm{V} \leftarrow+$ Infinity $z x \_$flag $\leftarrow 1$ | $\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)$ | $v \leftarrow+$ Zero | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{Q}(\text { srcc }) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ |
| +Infinity | $\begin{aligned} & \mathrm{V} \leftarrow d Q N a N \\ & \text { vxidi_flag } \leftarrow 1 \end{aligned}$ | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow+$ Infinity | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vxidi_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{aligned} & \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src2 }) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ |
| QNaN | $\mathrm{V} \leftarrow \mathrm{SrC1}$ | $\mathrm{V} \leftarrow \mathrm{SrCl}$ | $\mathrm{V} \leftarrow \mathrm{SrCl}$ | $\mathrm{V} \leftarrow \mathrm{SrC1}$ | $\mathrm{V} \leftarrow \mathrm{SrCl}$ | $\mathrm{V} \leftarrow \mathrm{SrCl}$ | $\mathrm{V} \leftarrow \mathrm{SrC1}$ | $\begin{aligned} & \mathrm{V} \leftarrow \text { src1 } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| SNaN | $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{Q} \text { (src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{array}{\|l\|} \hline v \leftarrow Q(s r c 1) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ | $\begin{array}{\|l\|} \hline v \leftarrow Q(\text { src1 } 1) \\ \text { vxsnan_flag } \leftarrow 1 \end{array}$ | $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{array}{\|l\|} \hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ \mathrm{vxsnan} \text { _lag } \leftarrow 1 \end{array}$ | $\begin{aligned} & \hline \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{array}{\|l\|} \hline V \leftarrow Q(\text { src } 1) \\ v x s n a n \_f l a g ~ \end{array} 1$ | $\begin{array}{\|l\|} \hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ \text { vxsnan_flag } \leftarrow 1 \end{array}$ |
| Explanation: |  |  |  |  |  |  |  |  |
| src1 | The double-precision floating-point value in doubleword element $i$ of VSR[XA] (where $i \in\{0,1\}$ ). |  |  |  |  |  |  |  |
| src2 | The double-precision floating-point value in doubleword element $i$ of VSR[XB] (where $i \in\{0,1\}$ ). |  |  |  |  |  |  |  |
| dQNaN | Default quiet NaN (0x7FF8_0000_0000_0000). |  |  |  |  |  |  |  |
| NZF | Nonzero finite number. |  |  |  |  |  |  |  |
| Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). |  |  |  |  |  |  |  |
| $D(x, y)$ | Return the normalized quotient of floating-point value x divided by floating-point value y , having unbounded range and precision. Return a QNaN with the payload of x . |  |  |  |  |  |  |  |
| $Q(x)$ |  |  |  |  |  |  |  |  |
| $v$ | The intermediate result having unbounded signficand precision and unbounded exponent range. |  |  |  |  |  |  |  |

Table 118.Actions for xvdivdp (element i)

VSX Vector Divide Single-Precision XX3-form
xvdivsp $\quad X T, X A, X B$


| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

$$
\text { do } \mathrm{i}=0 \text { to } 127 \text { by } 32
$$

reset_xflags()
$\operatorname{src} 1 \quad \leftarrow \operatorname{VSR}[X A]\{i: i+31\}$
$\operatorname{src} 2 \quad \leftarrow \operatorname{VSR}[X B]\{i: i+31\}$
v\{0:inf $\} \quad \leftarrow$ DivideSP(src1,src2)
result $\{\mathrm{i}: i+31\} \leftarrow$ RoundToSP(RN, v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxidi_flag) then SetFX(VXIDI)
if(vxisi_flag) then SetFX(VXZDZ)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
if(zx_flag) then $\operatorname{SetFX}(Z X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vxidi_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vxzdz_flag)
ex_flag $\leftarrow$ ex_flag | (OE \& ox_flag)
ex_flag $\leftarrow$ ex_flag | (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (ZE \& zx_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
end
if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3, do the following. Let src1 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X A]$.

Let src2 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.
src1 is divided ${ }^{[1]}$ by src2, producing a quotient having unbounded range and precision.

The quotient is normalized ${ }^{[2]}$.
See Table 119.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

[^3]|  | src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| -Infinity | $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{dQNaN} \\ & \text { vxidi_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{V} \leftarrow-\operatorname{lnfinity}$ | $V \leftarrow-$ Infinity | $\mathrm{V} \leftarrow-\operatorname{lnfinity}$ | $V \leftarrow-$ Infinity | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vxidi_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{v} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l\|} \hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 2) \\ \text { vxsnan_flag } \leftarrow 1 \end{array}$ |
| -NZF | $\mathrm{V} \leftarrow+$ Zero | $\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)$ | $\mathrm{V} \leftarrow+$ Infinity Zx_flag $\leftarrow 1$ | $\begin{aligned} & \mathrm{V} \leftarrow-\text { Infinity } \\ & \mathrm{zx} \text { flag } \leftarrow 1 \end{aligned}$ | $\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)$ | $v \leftarrow$-Zero | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & v \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| -Zero | $\mathrm{V} \leftarrow+$ Zero | $v \leftarrow+$ Zero | $\begin{aligned} & \hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vxzdz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \mathrm{vxzdz} \text { _lag } \leftarrow 1 \end{aligned}$ | $\mathrm{V} \leftarrow$-Zero | $v \leftarrow$-Zero | $\mathrm{v} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \hline v \leftarrow Q(\text { src2 }) \\ & v x \text { nnan_flag } \leftarrow 1 \end{aligned}$ |
| $\overline{\text { - }}$ +Zero | $\mathrm{V} \leftarrow$-Zero | $v \leftarrow$-Zero | $\begin{array}{\|l\|} \hline v \leftarrow d Q N a N \\ \text { vxzdz_lag } \leftarrow 1 \\ \hline \end{array}$ | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \mathrm{vxzdz} \mathrm{\_} \mathrm{flag} \mathrm{\leftarrow 1} \end{aligned}$ | $v \leftarrow+$ Zero | $v \leftarrow+$ Zero | $\mathrm{v} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & v \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| ¢ +NZF | $\mathrm{V} \leftarrow$-Zero | $\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)$ |  | $\mathrm{V} \leftarrow+$ Infinity zx_flag $\leftarrow 1$ | $\mathrm{V} \leftarrow \mathrm{D}(\mathrm{src} 1, \mathrm{src} 2)$ | $v \leftarrow+$ Zero | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l} \hline v \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| +Infinity | $\begin{aligned} & \hline v \leftarrow d Q N a N \\ & \text { vxidi_flag } \leftarrow 1 \end{aligned}$ | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $\mathrm{v} \leftarrow \mathrm{dQNaN}$ vxidi_flag $\leftarrow 1$ | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l\|} \hline v \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \end{array}$ |
| QNaN | $v \leftarrow \operatorname{src} 1$ | $\mathrm{V} \leftarrow \mathrm{SrC1}$ | $\mathrm{V} \leftarrow \mathrm{SrCl}$ | $\mathrm{v} \leftarrow \mathrm{SrCl}$ | $\mathrm{V} \leftarrow \mathrm{SrC1}$ | $\mathrm{V} \leftarrow \mathrm{SrCl}$ | $\mathrm{v} \leftarrow \mathrm{SrC1}$ | $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{src} 1 \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| SNaN | $\begin{aligned} & \hline \mathrm{V} \leftarrow \mathrm{Q}(\text { src1 } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src1 } 1) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline v \leftarrow Q(\text { src1 } 1) \\ & v x s n a n \_f l a g \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src1 } 1) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src1 } 1) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ |
| Explanation: |  |  |  |  |  |  |  |  |
| src1 | The single-precision floating-point value in word element $i$ of VSR[XA] (where $i \in\{0,1,2,3\}$ ). |  |  |  |  |  |  |  |
| src2 | The single-precision floating-point value in word element $i$ of VSR[XB] (where $i \in\{0,1,2,3\}$ ). |  |  |  |  |  |  |  |
| dQNaN | Default quiet NaN (0x7FC0_0000). |  |  |  |  |  |  |  |
| NZF | Nonzero finite number. |  |  |  |  |  |  |  |
| Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). |  |  |  |  |  |  |  |
| $D(x, y)$ | Return the normalized quotient of floating-point value $x$ divided by floating-point value $y$, having unbounded range and precision. Note: If $x=-y$, $v$ is considered to be an exact-zero-difference result (Rezd). |  |  |  |  |  |  |  |
| $Q(x)$ | Return a QNaN with the payload of $x$. |  |  |  |  |  |  |  |
| $v$ | The intermediate result having unbounded signficand precision and unbounded exponent range. |  |  |  |  |  |  |  |

Table 119.Actions for xvdivsp (element i)

VSX Vector Insert Exponent Double-Precision XX3-form

| Xviexpdp |
| :--- |
| 60  T  A  B  248 <br> 0  6  11  16  21 |

if MSR. VSX=0 then VSX_Unavailable(l)
do $\mathrm{i}=0$ to 1
srcl $\leftarrow \operatorname{VSR}[32 x A X+A]$. dword[i]
$\operatorname{src} 2 \leftarrow \operatorname{VSR}[32 \times B X+B]$. dword $[i]$
$\operatorname{VSR}[32 \times T X+T]$.dword[i].bit[0] $\leftarrow$ srcl.bit[0]
VSR[32xTX + T] dword[i], bit[1:11] $\leftarrow \operatorname{src} 2$. bit[53:63]
VSR[32xTX+T], dword[i], bit [12:63] $\operatorname{sic} 1$, bit [12:63]
end
Let $X T$ be the sum $32 \times T X+T$.
Let $X A$ be the sum $32 \times A X+A$.
Let $X B$ be the sum $32 \times B X+B$.
For each integer value i from 0 to 1, do the following. Let srcl be the unsigned integer value in doubleword element $i$ of VSR[ XA].

Let srcz be the unsigned integer value in doubleword element $i$ of VSR[XB].

The contents of bits 0 of $\operatorname{src} 1$ are placed into bit 0 of doubleword element $i$ of VSR[ XT].

The contents of bits 53:63 of src2 are placed into bits 1:11 of doubleword element $i$ of VSR[XT].

The contents of bits $12: 63$ of srcl are placed into bits 12:63 of doubleword element $i$ of VSR[ XT].

## Special Registers Altered:

None

VSX Vector Insert Exponent Single-Precision XX3-form

if MSR.VSX=0 then VSX_Unavailable()
do $i=0$ to 3
$\operatorname{srcl} \leftarrow \operatorname{VSR}[32 x A X+A]$, word $[i]$
$\operatorname{srcL} \leftarrow \operatorname{VSR}[32 \times B X+B]$. word $[i]$
VSR[32 $x T X+T]$, word[i].bit[0] $\leftarrow \operatorname{srclabit[0]}$
VSR[ $32 \times T X+T]$, word[i], bit[1:8] $\leftarrow$ src2. bit[24:31]
$\operatorname{VSR}[32 \times T X+T]$, word[i], bit $[9: 31] \leftarrow \operatorname{srcl}$, bit [9:31]
end
Let $X T$ be the sum $32 \times T X+T$.
Let $X A$ be the sum $32 \times A X+A$.
Let $X B$ be the sum $32 \times B X+B$.

For each integer value i from 0 to 3 , do the following. Let srcl be the unsigned integer value in word element $i$ of $V S R[X A]$.

Let $\operatorname{src} 2$ be the unsigned integer value in word element i of VSR[XB].

The contents of bits 0 of $\operatorname{srcl}$ are placed into bit 0 of word element $i$ of $V S R[X T]$.

The contents of bits 24:31 of $\operatorname{src}$ 2 are placed into bits $1: 8$ of word element $i$ of VSR[XT].

The contents of bits $9: 31$ of $\operatorname{srcl}$ are placed into bits 9:31 of word element $i$ of VSR[ XT] .

## Special Registers Altered:

None

VSR Data Layout for xviexpdp

| srcl | VSR[ XA$]. \mathrm{dword}$ [0] | VSR[ XA$]. \mathrm{dword}$ [1] |  |
| :---: | :---: | :---: | :---: |
| srct | VSR[ XB]. dword[0] | VSR[ XB$]. \mathrm{dword}$ [1] |  |
| tgt | VSR[ XT] , dword[0] | VSR[ XT$]$. dword[1] |  |
|  |  |  | 127 |

## VSR Data Layout for xviexpsp

| srcl | VSR[ XA], word[0] | VSR[ XA]. word[ 1] | VSR[ XA], word[ 2] | VSR[ XA]. word[3] |
| :---: | :---: | :---: | :---: | :---: |
| sic2 | VSR[ XB]. word[0] | VSR[ XB]. word [ 1] | VSR[ XB]. word[ 2] | VSR[ XB]. word[3] |
| tgt | VSR[ XT]. word[0] | VSR[ XT]. word[ 1] | VSR[ XT]. word[ 2] | VSR[ XT] . word[3] |
|  | 32 |  | 64 | 96 |

## VSX Vector Multiply-Add Double-Precision XX3-form

$$
\text { xvmaddadp } \quad X T, X A, X B
$$



| 60 | 6 | T |  |  |  |  | 21 | 105 | \|AxBx|TX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |


| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0$ 0b0 |

do $i=0$ to 127 by 64
reset_xflags()
$\operatorname{src1} \leftarrow \operatorname{VSR}[X A]\{i: i+63\}$
$\operatorname{src} 2 \leftarrow$ "xvmaddadp" ? VSR[XT]\{i:i+63\} : VSR[XB]\{i:i+63\}
src3 $\leftarrow$ "xvmaddadp" ? VSR[XB]\{i:i+63\} : VSR[XT]\{i:i+63\}
v\{0:inf\} $\leftarrow$ MultiplyAddDP(src1,src3,src2)
result $\{\mathrm{i}: i+63\} \leftarrow$ RoundToDP(RN, v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vximz_flag) then SetFX(VXIMZ)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(xx_flag) then $\operatorname{SetFX}(X X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vximz_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vxisi_flag)
ex_flag $\leftarrow$ ex_flag | (OE \& ox_flag)
ex_flag $\leftarrow$ ex_flag | (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
end
if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following.
For $\boldsymbol{x v m a d d a d p , ~ d o ~ t h e ~ f o l l o w i n g . ~}$

- Let src1 be the double-precision floating-point operand in doubleword element i of VSR[XA].
- Let src2 be the double-precision floating-point operand in doubleword element i of VSR[XT].
- Let src3 be the double-precision floating-point operand in doubleword element i of VSR[XB].

For xvmaddmdp, do the following.

- Let src1 be the double-precision floating-point operand in doubleword element i of VSR[XA].
- Let src2 be the double-precision floating-point operand in doubleword element i of VSR[XB].
- Let src3 be the double-precision floating-point operand in doubleword element i of VSR[XT].
src1 is multiplied ${ }^{[1]}$ by src3, producing a product having unbounded range and precision.

See part 1 of Table 120.
src2 is added ${ }^{[2]}$ to the product, producing a sum having unbounded range and precision.

The sum is normalized ${ }^{[3]}$.
See part 2 of Table 120.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element i of VSR[XT] in double-precision format.

See Table 106, "Vector Floating-Point Final Result," on page 663.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

Special Registers Altered<br>FX OX UX XX VXSNAN VXISI VXIMZ

## Version 3.0




| src3 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow-$ Infinity | $p \leftarrow-$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow+$ Infinity | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow+$ Zero | $p \leftarrow-$ Zero | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow+$ Zero | $p \leftarrow+$ Zero | $p \leftarrow-$ Zero | $p \leftarrow-$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow$ Src3 | $\begin{aligned} & p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow-$ Zero | $p \leftarrow-$ Zero | $p \leftarrow+$ Zero | $p \leftarrow+$ Zero | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow-$ Zero | $p \leftarrow+$ Zero | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow+$ Infinity | $\begin{array}{\|l\|} \hline p \leftarrow d Q N a N \\ \text { vximz_flag } \leftarrow 1 \\ \hline \end{array}$ | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{array}{\|l} \hline p \leftarrow Q(\text { src } 3) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{srcl}$ | $p \leftarrow \operatorname{src} 1$ | $\begin{aligned} & p \leftarrow \operatorname{src} 1 \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & p \leftarrow Q(\text { src } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & P \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |



| src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\mathrm{V} \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $\mathrm{V} \leftarrow-$ Infinity | $\begin{array}{\|l\|} \hline \mathrm{V} \leftarrow \mathrm{~d} Q \mathrm{NaN} \\ \text { vxisi_flag } \leftarrow 1 \\ \hline \end{array}$ | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{array}{\|l\|} \hline \mathrm{V} \leftarrow Q \text { (src2) } \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $v \leftarrow A(p, s r c 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{V} \leftarrow \mathrm{A}(\mathrm{p}, \mathrm{src} 2)$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{array}{\|l} \hline v \leftarrow Q \text { (src2) } \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\mathrm{V} \leftarrow$-Zero | $v \leftarrow$ Rezd | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{array}{\|l} \hline v \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \end{array}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $v \leftarrow$ Rezd | $\mathrm{V} \leftarrow+$ Zero | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{Src} 2$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & v x \text { vnan_flag } \leftarrow 1 \end{aligned}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{A}(\mathrm{p}, \mathrm{src} 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{V} \leftarrow \mathrm{A}(\mathrm{p}, \mathrm{src} 2)$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{array}{\|l\|} \hline v \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \end{array}$ |
| $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{~d} Q \mathrm{NaN} \\ & \text { vxisi_flag } \leftarrow 1 \end{aligned}$ | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & v x \text { nnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\begin{aligned} & v \leftarrow p \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $V \leftarrow p$ | $v \leftarrow p$ | $\mathrm{V} \leftarrow \mathrm{Src} 2$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |


| Explanation: |  |
| :---: | :---: |
| src1 | The double-precision floating-point value in doubleword element $i$ of VSR[XA] (where $i \in\{0,1\}$ ). |
| src2 | For $\boldsymbol{x} v$ maddadp, the double-precision floating-point value in doubleword element i of $\operatorname{VSR}[\mathrm{XT}]$ (where $\mathrm{i} \in\{0,1\}$ ). For $\boldsymbol{x v m a d d m d p}$, the double-precision floating-point value in doubleword element $i$ of VSR[XB] (where $i \in\{0,1\}$ ). |
| src3 | For $\boldsymbol{x} v m a d d a d p$, the double-precision floating-point value in doubleword element i of $\operatorname{VSR}[\mathrm{XB}]$ (where $\mathrm{i} \in\{0,1\}$ ). For xvmaddmdp, the double-precision floating-point value in doubleword element i of VSR[XT] (where i $\in\{0,1\}$ ). |
| dQNaN | Default quiet NaN (0x7FF8_0000_0000_0000). |
| NZF | Nonzero finite number. |
| Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| A(x,y) | Return the normalized sum of floating-point value $x$ and floating-point value $y$, having unbounded range and precision. Note: If $x=-y, v$ is considered to be an exact-zero-difference result (Rezd). |
| $\mathrm{M}(\mathrm{x}, \mathrm{y})$ | Return the normalized product of floating-point value $x$ and floating-point value $y$, having unbounded range and precision. |
| p | The intermediate product having unbounded range and precision. |
| v | The intermediate result having unbounded range and precision. |

Table 120.Actions for xvmadd(alm)dp

VSX Vector Multiply-Add Single-Precision XX3-form
xvmaddasp $\quad X T, X A, X B$

| 06 | 6 | 11 | A | 16 | B | 21 | 65 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| xvmaddmsp |  | XT, XA, XB |  |  |  |  |  |  |



| $X T$ | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

do $\mathrm{i}=0$ to 127 by 32
reset_xflags()
$\operatorname{src1} \leftarrow \operatorname{VSR}[\mathrm{XA}]\{i: 1+31\}$
$\operatorname{src} 2 \leftarrow{ }^{*}$ "vmaddasp" ? VSR[XT]\{i:i+31\}: VSR[XB]\{i:i+31\}
src3 $\leftarrow$ "xvmaddasp" ? VSR[XB]\{i:i+31\}: VSR[XT]\{i:i+31\}
v\{0:inf\} $\leftarrow$ MultiplyAddSP(src1, src3,src2)
result $\{i: i+63\} \leftarrow \operatorname{RoundToSP}(R N, v)$
if(vxsnan_flag) then SetFX(vXSNAN)
if(vximz_flag) then SetFX(VXIMZ)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then $\operatorname{SetFX}(0 X)$
if(ux_flag) then SetFX(UX)
if(xx_flag) then $\operatorname{SetFX}(x X)$
ex_flag $\leftarrow$ ex_flag \| (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag \| (VE \& vximz_flag)
ex_flag $\leftarrow$ ex_flag \| (VE \& vxisi_flag)
ex_flag $\leftarrow$ ex_flag | (OE \& ox_flag)
ex_flag $\leftarrow$ ex_flag | (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
end

$$
\text { if( ex_flag = } 0 \text { ) then VSR }[X T] \leftarrow \text { result }
$$

Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3 , do the following.
For xvmaddasp, do the following.

- Let src1 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X A]$.
- Let src2 be the single-precision floating-point operand in word element i of VSR[XT].
- Let src3 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.

For xvmaddmsp, do the following.

- Let src1 be the single-precision floating-point operand in word element i of VSR[XA].
- Let src2 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.
- Let src3 be the single-precision floating-point operand in word element i of VSR[XT].
src1 is multiplied ${ }^{[1]}$ by src3, producing a product having unbounded range and precision.

See part 1 of Table 121.
src2 is added ${ }^{[2]}$ to the product, producing a sum having unbounded range and precision.

The sum is normalized ${ }^{[3]}$.
See part 2 of Table 121.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into word element $i$ of $\mathrm{VSR}[\mathrm{XT}]$ in single-precision format.

See Table 106, "Vector Floating-Point Final Result," on page 663.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

FX OX UX XX VXSNAN VXISI VXIMZ

[^4]VSR Data Layout for xvmadd(alm)sp
src1 = VSR[XA]

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

src2 = xsmaddasp ? VSR[XT] : VSR[XB]

| SP | SP | SP | SP |
| :---: | :---: | :---: | :---: |
| Src3 = xsmaddasp ? VSR[XB] : VSR[XT] |  |  |  |
| SP | SP | SP | SP |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

| SP | SP | SP |  | SP |  |
| :--- | :--- | :--- | :--- | :--- | :---: |
| 0 | 32 |  | 64 |  |  |



| src3 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow-$ Infinity | $p \leftarrow-$ Infinity | $p \leftarrow$ src3 | $\begin{aligned} & p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow+$ Infinity | $p \leftarrow M(s r c 1, s r c 3)$ | $p \leftarrow+$ Zero | $p \leftarrow-$ Zero | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow$ src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow+$ Zero | $\mathrm{p} \leftarrow+$ Zero | $\mathrm{p} \leftarrow-$ Zero | $p \leftarrow-$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{p} \leftarrow$-Zero | $\mathrm{p} \leftarrow$-Zero | $\mathrm{p} \leftarrow+$ Zero | $p \leftarrow+$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow M(s r c 1, s r c 3)$ | $\mathrm{p} \leftarrow-$ Zero | $\mathrm{p} \leftarrow+$ Zero | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow$ src3 | $\begin{aligned} & \mathrm{p} \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \mathrm{p} \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $\begin{aligned} & p \leftarrow \text { src1 } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & \hline p \leftarrow Q(\text { src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{P} \leftarrow Q(\text { src } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{p} \leftarrow Q(\text { src } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{p} \leftarrow Q(\text { src } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |



| src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\mathrm{V} \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $\begin{array}{\|l\|} \hline v \leftarrow d Q N a N \\ \text { vxisi_flag } \leftarrow 1 \end{array}$ | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \hline v \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $v \leftarrow A(p, s r c 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow A\left(p, s c^{\prime} 2\right)$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{v} \leftarrow \mathrm{SrC2}$ | $\mathrm{V} \leftarrow$-Zero | $v \leftarrow$ Rezd | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{array}{\|l\|} \hline v \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $v \leftarrow$ Rezd | $\mathrm{v} \leftarrow+$ Zero | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $v \leftarrow A(p, s r c 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow A(p, s r c 2)$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \hline v \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{dQNaN} \\ & \text { vxisi_flag } \leftarrow 1 \end{aligned}$ | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\begin{aligned} & v \leftarrow p \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |

## Explanation:

| src1 | The single-precision floating-point value in word element i of VSR[XA] (where $i \in\{0,1,2,3\}$ ). |
| :---: | :---: |
| src2 | For $\boldsymbol{x v m a d d a s p}$, the single-precision floating-point value in word element $i$ of VSR[XT] (where $i \in\{0,1,2,3\}$ ). For $\boldsymbol{x v m a d d m s p}$, the single-precision floating-point value in word element $i$ of VSR[XB] (where $i \in\{0,1,2,3\}$ ). |
| src3 | For $\boldsymbol{x v m a d d a s p}$, the single-precision floating-point value in word element $i$ of VSR[XB] (where $i \in\{0,1,2,3\}$ ). For xvmaddmsp, the single-precision floating-point value in word element $i$ of VSR[XT] (where $i \in\{0,1,2,3\}$ ). |
| dQNaN | Default quiet NaN (0x7FC0_0000). |
| NZF | Nonzero finite number. |
| Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| A(x,y) | Return the normalized sum of floating-point value $x$ and floating-point value $y$, having unbounded range and precision. Note: If $x=-y, v$ is considered to be an exact-zero-difference result (Rezd). |
| M ( $\mathrm{x}, \mathrm{y}$ ) | Return the normalized product of floating-point value $x$ and floating-point value y , having unbounded range and precision. |
| p | The intermediate product having unbounded range and precision. |
| v | The intermediate result having unbounded range and precision. |

Table 121.Actions for xvmadd(alm)sp

## VSX Vector Maximum Double-Precision XX3-form


$i=0$ to 127 by 64
reset_xflags()
SrC
result $\{i: i+63\} \leftarrow \operatorname{MaximumDP}(\operatorname{src} 1, s r c 2)$
if(vxsnan_flag) then SetFX(VXSNAN)
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result

Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let src1 be the double-precision floating-point operand in doubleword element $i$ of $\operatorname{VSR}[X A]$.

Let src2 be the double-precision floating-point operand in doubleword element $i$ of VSR[XB].

If src1 is greater than src2, src1 is placed into doubleword element $i$ of VSR[XT] in double-precision format. Otherwise, src2 is placed into doubleword element i of $\operatorname{VSR}[\mathrm{XT}]$ in double-precision format.

The maximum of +0 and -0 is +0 . The maximum of a QNaN and any value is that value. The maximum of any value and an SNaN when VE=0 is that SNaN converted to a QNaN .

See Table 122.
If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

Special Registers Altered
FX VXSNAN

VSR Data Layout for xvmaxdp
src1 = VSR[XA]

| DP | DP |
| :---: | :---: |

$\operatorname{src} 2=\mathrm{VSR}[\mathrm{XB}]$

| DP | DP |
| :---: | :---: |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

| DP | DP |
| :--- | :--- |
| 0 | 64 |


|  | src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| -Infinity | T(sicl) | T(sro2) | T(sro2) | T(sro2) | T (scr2) | T (scre) | T(scri) | $T(Q(s \mathrm{sc} 2))$ fx(VXSNAN) |
| -NZF | T (sicl) | T(M(ssci, ,sc2)) | T (sro2) | T (sro2) | T (scr2) | T (scre2) | T(scri) | $T\left(Q\left(s r^{2}\right)\right)$ fx(VXSNAN) |
| -Zero | T(srci) | T (stri) | T (scri) | T(str2) | T(str2) | $T($ scor $)$ | T (scri) | $T\left(Q\left(s r^{2}\right)\right)$ fx(VXSNAN) |
| - +Zero | T(srci) | T (stri) | T(scri) | T (scol) | T(str2) | $T(s \mathrm{c} 2)$ | T (scri) | $T\left(Q\left(s r^{2}\right)\right)$ tx(VXSNAN) |
| +NZF | T(srci) | T(srci) | T (stri) | T(srci) | T(MMssc1,sc2)) | T(str2) | T(scri) | $T\left(Q\left(s r_{2}\right)\right)$ tx(VXSNAN) |
| +Infinity | T(srci) | T (stri) | T (stc1) | T (stc1) | T (stc1) | T(stri) | T (scol) | $T\left(Q\left(s s_{0}\right)\right)$ fx(VXSNAN) |
| QNaN | T(sro2) | T(str2) | T(str2) | T(sroz) | T(str2) | T(str2) | T(ssc1) | T (src1) fx(VXSNAN) |
| SNaN | $T(Q(\operatorname{sc} 11))$ fxVXSNAN | $T(Q(s \mathrm{c} 11))$ fx(VXSNAN) | $T\left(Q\left(s c^{\prime} 1\right)\right)$ txVXSNAN | $T(Q(s \mathrm{cr} 1))$ tx(XXSNAN | $T(Q($ scicl $))$ fx(VXSNAN) | $T(Q($ srcil $))$ fx(VXSNAN) | T(Q(scil)) tx(VXSNAN) | $T(Q(s c 1))$ fx(XXSNAN) |


| Explanation: |  |
| :---: | :---: |
| src1 | The double-precision floating-point value in doubleword element $i$ of VSR[XA] (where $i \in\{0,1\}$ ). |
| src2 | The double-precision floating-point value in doubleword element i of VSR[XT] (where $\mathrm{i} \in\{0,1\}$ ). |
| NZF | Nonzero finite number. |
| Q(x) | Return a QNaN with the payload of x . |
| M ( $\mathrm{x}, \mathrm{y}$ ) | Return the greater of floating-point value $x$ and floating-point value y . |
| T(x) | The value $x$ is placed in doubleword element $i(i \in\{0,1\})$ of VSR[XT] in double-precision format. FPRF, FR and FI are not modified. |
| fx(x) | If $x$ is equal to $0, F X$ is set to $1 . x$ is set to 1 . |
| VXSNAN | Floating-point Invalid Operation Exception (SNaN). If VE=1, update of VSR[XT] is suppressed. |

Table 122.Actions for xvmaxdp

VSX Vector Maximum Single-Precision XX3-form

| xvmaxsp |  | XT, XA, XB |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll}  & 60 \\ 0 & \end{array}$ | ${ }_{6} \quad \mathrm{~T}$ | ${ }_{11} A$ | 16 B | 2192 | $\begin{array}{l\|l\|l\|} \mathrm{AX} & \mathrm{BX} \\ 29 & 30 & \mathrm{TX} \\ \hline 1 \end{array}$ |
| XT $\quad \leftarrow T X \\|$ T |  |  |  |  |  |
| $X A \quad \leftarrow A X \\| A$ |  |  |  |  |  |
|  |  |  |  |  |  |
| ex_flag $\leftarrow 0 \mathrm{bb}$ |  |  |  |  |  |
| do i=0 to 127 by 32 |  |  |  |  |  |
| reset_xflags() |  |  |  |  |  |
| src1 $\leftarrow \operatorname{VSR}[\mathrm{XA}]\{\mathrm{i}: i+31\}$ |  |  |  |  |  |
| src2 $\leftarrow \operatorname{VSR}[\mathrm{XB}]\{\mathrm{i}: i+31\}$ |  |  |  |  |  |
| result $\{\mathrm{i}: 1+63\} \leftarrow$ MaximumSP(src1, src2) |  |  |  |  |  |
| if(vxsnan_flag) then SetFX(VXSNAN) |  |  |  |  |  |
| ex_flag $\leftarrow$ ex_flag \| (VE \& vxsnan_flag) |  |  |  |  |  |
| end |  |  |  |  |  |
| if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result |  |  |  |  |  |

VSR Data Layout for xvmaxsp
src1 = VSR[XA]

| SP | SP | SP | SP |
| :---: | :---: | :---: | :---: |
| Src2 $=$ VSR[XB] |  |  |  |
| SP | SP | SP | SP |

tgt $=\mathrm{VSR}[\mathrm{XT}]$


Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3, do the following. Let src1 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[\mathrm{XA}]$.

Let src2 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.

If src1 is greater than src2, src1 is placed into word element i of VSR[XT] in single-precision format. Otherwise, src2 is placed into word element i of VSR[XT] in single-precision format.

The maximum of +0 and -0 is +0 . The maximum of a QNaN and any value is that value. The maximum of any value and an SNaN when VE=0 is that SNaN converted to a QNaN.

See Table 123.
If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

FX VXSNAN

|  | src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| -Infinity | T(sicl) | T(sro2) | T(sro2) | T(sro2) | T (scr2) | T (scre) | T(scri) | $T(Q(s \mathrm{sc} 2))$ fx(VXSNAN) |
| -NZF | T (sicl) | T(M(ssci, ,sc2)) | T (sro2) | T (sro2) | T (scr2) | T (scre2) | T(scri) | $T\left(Q\left(s r^{2}\right)\right)$ fx(VXSNAN) |
| -Zero | T(srci) | T (stri) | T (scri) | T(str2) | T(str2) | $T($ scor $)$ | T (scri) | $T\left(Q\left(s r^{2}\right)\right)$ fx(VXSNAN) |
| - +Zero | T(srci) | T (stri) | T(scri) | T (scol) | T(str2) | $T(s \mathrm{c} 2)$ | T (scri) | $T\left(Q\left(s r^{2}\right)\right)$ tx(VXSNAN) |
| +NZF | T(srci) | T(srci) | T (stri) | T(srci) | T(MMssc1,sc2)) | T(str2) | T(scri) | $T\left(Q\left(s r_{2}\right)\right)$ tx(VXSNAN) |
| +Infinity | T(srci) | T (stri) | T (stc1) | T (stc1) | T (stc1) | T(stri) | T (scol) | $T\left(Q\left(s s_{0}\right)\right)$ fx(VXSNAN) |
| QNaN | T(sro2) | T(str2) | T(str2) | T(sroz) | T(str2) | T(str2) | T(ssc1) | T (src1) fx(VXSNAN) |
| SNaN | $T(Q(\operatorname{sc} 11))$ fxVXSNAN | $T(Q(s \mathrm{c} 11))$ fx(VXSNAN) | $T\left(Q\left(s c^{\prime} 1\right)\right)$ txVXSNAN | $T(Q(s \mathrm{cr} 1))$ tx(XXSNAN | $T(Q($ scicl $))$ fx(VXSNAN) | $T(Q($ srcil $))$ fx(VXSNAN) | T(Q(scil)) tx(VXSNAN) | $T(Q(s c 1))$ fx(XXSNAN) |


| Explanation: |  |
| :---: | :---: |
| src1 | The single-precision floating-point value in word element i of VSR[XA] (where $i \in\{0,1,2,3\}$ ). |
| src2 | The single-precision floating-point value in word element i of VSR[XT] (where $i \in\{0,1,2,3\}$ ). |
| NZF | Nonzero finite number. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| M ( $\mathrm{x}, \mathrm{y}$ ) | Return the greater of floating-point value x and floating-point value y . |
| T(x) | The value x is placed in word element $\mathrm{i}(\mathrm{i} \in\{0,1,2,3\})$ of VSR[XT] in single-precision format. FPRF, FR and FI are not modified. |
| $\mathrm{fx}(\mathrm{x})$ | If $x$ is equal to $0, F X$ is set to $1 . x$ is set to 1 . |
| VXSNAN | Floating-point Invalid Operation Exception (SNaN). If VE=1, update of VSR[XT] is suppressed. |

Table 123.Actions for xvmaxsp

## VSX Vector Minimum Double-Precision XX3-form



| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0$ b0 |

```
do i=0 to 127 by 64
    reset_xflags()
    src1 }\leftarrow\operatorname{VSR[XA]{i:i+63}
    src2 }\leftarrowVVR[XB]{i:i+63
    result{i:i+63} \leftarrow MinimumDP(src1,src2)
    if(vxsnan_flag) then SetFX(VXSNAN)
    ex_flag & ex_flag | (VE & vxsnan_flag)
end
```

if( ex_flag = 0 ) then VSR $[\mathrm{XT}] \leftarrow$ result

Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1, do the following. Let src1 be the double-precision floating-point operand in doubleword element $i$ of VSR[XA].

Let src2 be the double-precision floating-point operand in doubleword element $i$ of VSR[XB].

If src1 is less than src2, src1 is placed into doubleword element $i$ of VSR[XT] in double-precision format. Otherwise, src2 is placed into doubleword element i of $\operatorname{VSR}[\mathrm{XT}]$ in double-precision format.

The minimum of +0 and -0 is -0 . The minimum of a QNaN and any value is that value. The minimum of any value and an SNaN when $\mathrm{VE}=0$ is that SNaN converted to a QNaN.

See Table 124.
If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

Special Registers Altered
FX VXSNAN

VSR Data Layout for xvmindp
src1 = VSR[XA]

| DP | DP |
| :---: | :---: |
| src2 $=$ VSR[XB] |  |
| DP | DP |

$\operatorname{tgt}=\mathrm{VSR}[\mathrm{XT}]$

|  | DP |
| :--- | :--- |
| 0 | 64 |


|  | src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| -Infinity | T(sicl) | T(srci) | T(sric) | T(srci) | T (scri) | T(scri) | T(scri) | $T(Q(s \mathrm{sc} 2))$ fx(VXSNAN) |
| -NZF | T (sro2) | T(M(ssci, ,sc2)) | T (sic1) | T (scri) | T (scri) | T (scri) | T(scri) | $T\left(Q\left(s r^{2}\right)\right)$ fx(VXSNAN) |
| -Zero | T(sro2) | T(str2) | T (scri) | T (scri) | T (stri) | T(scri) | T (scri) | $T\left(Q\left(s r^{2}\right)\right)$ fx(VXSNAN) |
| - +Zero | T(sro2) | T(str2) | T (str2) | T (scol) | T (scol) | T(scri) | T (scri) | $T\left(Q\left(s r^{2}\right)\right)$ tx(VXSNAN) |
| +NZF | T(sroz) | T(sro2) | T(str2) | T(sro2) | T(MMssc1,sc2)) | T(scri) | T(scri) | $T\left(Q\left(s r_{2}\right)\right)$ tx(VXSNAN) |
| +Infinity | T(srce) | T(sro2) | T(str2) | T(sro2) | T(str2) | T(stri) | T (scol) | $T\left(Q\left(s s_{0}\right)\right)$ fx(VXSNAN) |
| QNaN | T(sro2) | T(str2) | T(str2) | T(sroz) | T(str2) | T(str2) | T(ssc1) | T (src1) fx(VXSNAN) |
| SNaN | $T(Q(\operatorname{sc} 11))$ fxVXSNAN | $T(Q(s \mathrm{c} 11))$ fx(VXSNAN) | $T\left(Q\left(s c^{\prime} 1\right)\right)$ txVXSNAN | $T(Q(s \mathrm{cr} 1))$ tx(XXSNAN | $T(Q($ scicl $))$ fx(VXSNAN) | $T(Q($ srcil $))$ fx(VXSNAN) | $T(Q(s \mathrm{cr} 1))$ fxVXSNAN | $T(Q(s c 1))$ fx(XXSNAN) |


| Explanation: |  |
| :---: | :---: |
| src1 | The double-precision floating-point value in doubleword element i of VSR[XA] (where $\mathrm{i} \in\{0,1\}$ ). |
| src2 | The double-precision floating-point value in doubleword element $i$ of VSR[XT] (where $i \in\{0,1\}$ ). |
| NZF | Nonzero finite number. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| M ( $\mathrm{x}, \mathrm{y}$ ) | Return the lesser of floating-point value $x$ and floating-point value y . |
| T(x) | The value $x$ is placed in doubleword element $i(i \in\{0,1\})$ of VSR[XT] in double-precision format. FPRF, FR and FI are not modified. |
| $\mathrm{fx}(\mathrm{x})$ | If $x$ is equal to $0, F X$ is set to $1 . x$ is set to 1 . |
| VXSNAN | Floating-point Invalid Operation Exception (SNaN). If VE=1, update of VSR[XT] is suppressed. |

Table 124.Actions for xvmindp

## VSX Vector Minimum Single-Precision XX3-form



VSR Data Layout for xvminsp
src1 = VSR[XA]

| SP | SP | SP | SP |
| :---: | :---: | :---: | :---: |
| Src2 $=$ VSR[XB] |  |  |  |
| SP | SP | SP | SP |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

| SP | SP |  | SP | SP |  |
| :--- | :--- | :--- | :--- | :--- | :---: |
| 0 | 32 | 64 |  | 96 |  |

Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3, do the following. Let src1 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[\mathrm{XA}]$.

Let src2 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.

If src1 is less than src2, src1 is placed into word element i of VSR[XT] in single-precision format. Otherwise, src2 is placed into word element i of VSR[XT] in single-precision format.

The minimum of +0 and -0 is -0 . The minimum of a QNaN and any value is that value. The minimum of any value and an SNaN when $\mathrm{VE}=0$ is that SNaN converted to a QNaN.

See Table 125.
If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

FX VXSNAN

|  | src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| -Infinity | T(sicl) | T(srci) | T(sric) | T(srci) | T (scri) | T(scri) | T(scri) | $T(Q(s \mathrm{sc} 2))$ fx(VXSNAN) |
| -NZF | T (sro2) | T(M(ssci, ,sc2)) | T (sic1) | T (scri) | T (scri) | T (scri) | T(scri) | $T\left(Q\left(s r^{2}\right)\right)$ fx(VXSNAN) |
| -Zero | T(sro2) | T(str2) | T (scri) | T (scri) | T (stri) | T(scri) | T (scri) | $T\left(Q\left(s r^{2}\right)\right)$ fx(VXSNAN) |
| - +Zero | T(sro2) | T(str2) | T (str2) | T (scol) | T (scol) | T(scri) | T (scri) | $T\left(Q\left(s r^{2}\right)\right)$ tx(VXSNAN) |
| +NZF | T(sroz) | T(sro2) | T(str2) | T(sro2) | T(MMssc1,sc2)) | T(scri) | T(scri) | $T\left(Q\left(s r_{2}\right)\right)$ tx(VXSNAN) |
| +Infinity | T(srce) | T(sro2) | T(str2) | T(sro2) | T(str2) | T(stri) | T (scol) | $T\left(Q\left(s s_{0}\right)\right)$ fx(VXSNAN) |
| QNaN | T(sro2) | T(str2) | T(str2) | T(sroz) | T(str2) | T(str2) | T(ssc1) | T (src1) fx(VXSNAN) |
| SNaN | $T(Q(\operatorname{sc} 11))$ fxVXSNAN | $T(Q(s \mathrm{c} 11))$ fx(VXSNAN) | $T\left(Q\left(s c^{\prime} 1\right)\right)$ txVXSNAN | $T(Q(s \mathrm{cr} 1))$ tx(XXSNAN | $T(Q($ scicl $))$ fx(VXSNAN) | $T(Q($ srcil $))$ fx(VXSNAN) | $T(Q(s \mathrm{cr} 1))$ fxVXSNAN | $T(Q(s c 1))$ fx(XXSNAN) |


| Explanation: |  |
| :---: | :---: |
| src1 | The single-precision floating-point value in word element i of VSR[XA] (where $i \in\{0,1,2,3\}$ ). |
| src2 | The single-precision floating-point value in word element i of VSR[XT] (where $i \in\{0,1,2,3\}$ ). |
| NZF | Nonzero finite number. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| M ( $\mathrm{x}, \mathrm{y}$ ) | Return the lesser of floating-point value $x$ and floating-point value y . |
| T(x) | The value x is placed in word element $\mathrm{i}(\mathrm{i} \in\{0,1,2,3\})$ of VSR[XT] in single-precision format. FPRF, FR and FI are not modified. |
| $\mathrm{fx}(\mathrm{x})$ | If $x$ is equal to $0, F X$ is set to $1 . x$ is set to 1 . |
| VXSNAN | Floating-point Invalid Operation Exception (SNaN). If VE=1, update of VSR[XT] is suppressed. |

Table 125.Actions for xvminsp

## VSX Vector Multiply-Subtract Double-Precision XX3-form

xvmsubadp $\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}$

| 60 | 6 |  | 11 | A |  | B | 21 | 113 | $\left\lvert\, \begin{aligned} & \text { AXPXXTX } \\ & 293031\end{aligned}\right.$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| xvmsubmdp |  |  |  |  |  |  |  |  |  |



| XT $\quad \leftarrow$ | $\leftarrow T X \\| T$ |
| :---: | :---: |
| $X A \quad \leftarrow$ | $\leftarrow A X \\| A$ |
| $X B \quad \leftarrow$ | $\leftarrow B X\|\mid B$ |
| ex_flag $\leftarrow$ | $\leftarrow 0 \mathrm{bO}$ |
| do i=0 to 127 by 64 |  |
| reset_xflags() |  |
| src1 $\leftarrow \operatorname{VSR}[\mathrm{XA}]\{\mathrm{i}: 1+63\}$ |  |
| src2 $\leftarrow$ "xvmsubadp" ? VSR[XT]\{i:i+63\}: VSR[XB]\{i:i+63\} |  |
| src3 $\leftarrow$ "xvmsubadp" ? VSR[XB]\{i:i+63\} : VSR[XT]\{i:i+63\} |  |
| v\{0:inf\} | $\} \quad \leftarrow$ MultiplyAddDP(src1,src3, NegateDP(src2)) |
| result $\{\mathrm{i}: 1+63\} \leftarrow$ RoundToDP(RN,v) |  |
| if(vxsnan_flag) then SetFX(VXSNAN) |  |
| if(vximz_flag) then SetFX(VXIMZ) |  |
| if(vxisi_flag) then SetFX(VXISI) |  |
| if(ox_flag) then SetFX(0X) |  |
| if(ux_flag) then SetFX(UX) |  |
| if(xx_flag) then SetFX(XX) |  |
| ex_flag | $\leftarrow$ ex_flag \| (VE \& vxsnan_flag) |
| ex_flag | $\leftarrow$ ex_flag \| (VE \& vximz_flag) |
| ex_flag | $\leftarrow$ ex_flag \| (VE \& vxisi_flag) |
| ex_flag | $\leftarrow$ ex_flag \| (OE \& 0x_flag) |
| ex_flag | $\leftarrow$ ex_flag \| (UE \& ux_flag) |
| ex_flag | $\leftarrow$ ex_flag \| (XE \& xx_flag) |
| nd |  |

if( ex_flag =0) then VSR $[\mathrm{XT}] \leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following.
For $\boldsymbol{x v m s u b a d p , ~ d o ~ t h e ~ f o l l o w i n g . ~}$

- Let src1 be the double-precision floating-point operand in doubleword element i of VSR[XA].
- Let src2 be the double-precision floating-point operand in doubleword element i of VSR[XT].
- Let src3 be the double-precision floating-point operand in doubleword element i of VSR[XB].

For xvmsubmdp, do the following.

- Let src1 be the double-precision floating-point operand in doubleword element i of VSR[XA].
- Let src2 be the double-precision floating-point operand in doubleword element i of VSR[XB].
- Let src3 be the double-precision floating-point operand in doubleword element i of VSR[XT].
src1 is multiplied ${ }^{[1]}$ by src3, producing a product having unbounded range and precision.

See part 1 of Table 126.
src2 is negated and added ${ }^{[2]}$ to the product, producing a sum having unbounded range and precision.

The sum is normalized ${ }^{[3]}$.
See part 2 of Table 126.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element i of VSR[XT] in double-precision format.

See Table 106, "Vector Floating-Point Final Result," on page 663.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

FX OX UX XX VXSNAN VXISI VXIMZ

## Version 3.0

| VSR Data Layout for xvmsub(alm)dp |
| :--- |
| src1 = VSR[XA] |
| DP DP <br> src2 = xvmsubadp ? VSR[XT] : VSR[XB]  <br> DP  <br> Drc3 = xvmsubadp ? VSR[XB] : VSR[XB]  <br> DP DP   <br> tgt = VSR[XT]    <br> DP    <br> 0    127 DP |



| src3 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow-$ Infinity | $p \leftarrow-$ Infinity | $p \leftarrow$ Src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $p \leftarrow+$ Infinity | $p \leftarrow M($ src1,src3 $)$ | $p \leftarrow+$ Zero | $\mathrm{p} \leftarrow$-Zero | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow$ src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow+$ Zero | $\mathrm{p} \leftarrow+$ Zero | $p \leftarrow-$ Zero | $p \leftarrow-$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow$ Src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow-$ Zero | $\mathrm{p} \leftarrow$-Zero | $p \leftarrow+$ Zero | $p \leftarrow+$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow$ src3 | $\begin{aligned} & p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow M($ src1,src3 $)$ | $p \leftarrow-$ Zero | $p \leftarrow+$ Zero | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow$ Src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $p \leftarrow$ src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\mathrm{p} \leftarrow \mathrm{SrC1}$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $\mathrm{p} \leftarrow \mathrm{src} 1$ | $\begin{aligned} & p \leftarrow \text { src1 } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & p \leftarrow Q(\text { src } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & P \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & P \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q \text { (src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |



| src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\mathrm{V} \leftarrow \mathrm{dQNaN}$ <br> vxisi_flag $\leftarrow 1$ | $\mathrm{v} \leftarrow-\operatorname{lnfinity}$ | $\mathrm{V} \leftarrow-$-nfinity | $v \leftarrow-$ lnfinity | $\mathrm{v} \leftarrow-$ lnfinity | $\mathrm{v} \leftarrow-$ Infinity | $\mathrm{v} \leqslant$ src2 | $\begin{aligned} & \hline \begin{array}{l} v \leftarrow Q(\text { src2 } 2) \\ v x s n a n \_l a g \leftarrow 1 \end{array} \end{aligned}$ |
| $\mathrm{V} \leftarrow+$ lnfinity | $v \leqslant S(p, s \mathrm{sc} 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{v} \leqslant \mathrm{S}(\mathrm{p}, \mathrm{src} 2)$ | $\mathrm{V} \leftarrow-$-nfinity | $\mathrm{v} \leqslant \mathrm{src} 2$ | $\begin{array}{\|l\|} \hline v \leftarrow Q(\text { srč2) } \\ v x s n a n \_l a g \leftarrow 1 \end{array}$ |
| $v \leftarrow+$ lnfinity | v - Src2 | $v \leftarrow-$ Zero | $v \leftarrow$ Rezd | $\mathrm{v} \leftarrow-\mathrm{Src} 2$ | $v \leftarrow-$ Infinity | $\mathrm{v} \leqslant$ src2 | $\begin{array}{\|l\|} \hline v \leftarrow Q(\text { srcc2 }) \\ \text { vxsnan_lag } \leftarrow 1 \\ \hline \end{array}$ |
| $\mathrm{V} \leftarrow+$ lnfinity | v ¢-SrC2 | $v \leftarrow$ Rezd | $v \leftarrow+$ Zero | $\mathrm{v} \leftarrow-\mathrm{Src} 2$ | $\mathrm{V} \leftarrow-$-nfinity | v ¢ SrC2 | $\begin{aligned} & \hline \mathrm{V} \leftarrow Q \text { (src2) } \\ & \mathrm{vxsnan} \text { _lag } \leftarrow 1 \end{aligned}$ |
| $\mathrm{V} \leftarrow+$ lnfinity | $v \leqslant S(p, s \mathrm{sc} 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{v} \leqslant \mathrm{S}(\mathrm{p}, \mathrm{src} 2)$ | $v \leftarrow-$ Infinity | v ¢ Src2 | $\begin{aligned} & \hline \mathrm{V} \leftarrow Q(\text { srcc2 }) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |
| $\mathrm{V} \leftarrow+$ lnfinity | $v \leftarrow+$ lnfinity | $v \leftarrow+$ lnfinity | $v \leftarrow+$ lnfinity | $v \leftarrow+$ lnfinity | $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{~d} \mathrm{NaN} \\ & \text { vxisiflag } \leftarrow 1 \end{aligned}$ | $\mathrm{v} \leqslant \mathrm{SrO} 2$ | $\begin{array}{\|l\|} \hline v \leftarrow Q(\text { srcc } 2) \\ v x s n a n \_l a g \leftarrow 1 \end{array}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\begin{array}{\|l\|} \hline v \leftarrow p \\ \text { vxsnan_flag } \leftarrow 1 \end{array}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | V ¢ SrC2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srč2) } \\ & \mathrm{vxsnn} \text { _flag } \leftarrow 1 \end{aligned}$ |


| Explanation: |  |
| :---: | :---: |
| src1 | The double-precision floating-point value in doubleword element $i$ of VSR[XA] (where $i \in\{0,1\}$ ). |
| src2 | For $\boldsymbol{x} v m s u b a d p$, the double-precision floating-point value in doubleword element i of $\operatorname{VSR}[\mathrm{XT}]$ (where $\mathrm{i} \in\{0,1\}$ ). For $\boldsymbol{x v m s u b m d p}$, the double-precision floating-point value in doubleword element $i$ of VSR[XB] (where $i \in\{0,1\}$ ). |
| src3 | For $\boldsymbol{x} v m s u b a d p$, the double-precision floating-point value in doubleword element i of $\operatorname{VSR}[\mathrm{XB}]$ (where $\mathrm{i} \in\{0,1\}$ ). For xvmsubmdp, the double-precision floating-point value in doubleword element $i$ of VSR[XT] (where $i \in\{0,1\}$ ). |
| dQNaN | Default quiet NaN (0x7FF8_0000_0000_0000). |
| NZF | Nonzero finite number. |
| Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| $\mathrm{S}(\mathrm{x}, \mathrm{y})$ | Return the normalized sum of floating-point value $x$ and negated floating-point value $y$, having unbounded range and precision. Note: If $x=y, v$ is considered to be an exact-zero-difference result (Rezd). |
| $\mathrm{M}(\mathrm{x}, \mathrm{y})$ | Return the normalized product of floating-point value $x$ and floating-point value $y$, having unbounded range and precision. |
| p | The intermediate product having unbounded range and precision. |
| $\checkmark$ | The intermediate result having unbounded range and precision. |

Table 126.Actions for xvmsub(alm)dp

VSX Vector Multiply-Subtract Single-Precision XX3-form
xvmsubasp $\quad X T, X A, X B$

| 06 | 6 | 11 | A | 16 | B | 21 | 81 | $\left\lvert\, \begin{aligned} & \text { AXXXXTX } \\ & 293031\end{aligned}\right.$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| xvmsubmsp |  | XT, XA, XB |  |  |  |  |  |  |



| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

do $\mathrm{i}=0$ to 127 by 32
reset_xflags()
$\operatorname{src} 1 \quad \leftarrow \operatorname{VSR}[X A]\{i: i+31\}$
src2 $\leftarrow$ "xvmsubasp" ? VSR[XT]\{i:i+31\}: VSR[XB]\{i:i+31\}
src3 $\leftarrow$ "xvmsubasp" ? VSR[XB]\{i:i+31\}: VSR[XT]\{i:i+31\}
v\{0:inf\} $\leftarrow$ MultiplyAddSP(src1,src3,NegateSP(src2))
result $\{\mathrm{i}: 1+31\} \leftarrow \operatorname{RoundToSP}(\mathrm{RN}, \mathrm{v})$
if(vxsnan_flag) then SetFX(vXSNAN)
if(vximz_flag) then SetFX(vXIMZ)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then $\operatorname{SetFX}(0 X)$
if(ux_flag) then SetFX(UX)
if(xx_flag) then $\operatorname{SetFX}(x X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag \| (VE \& vximz_flag)
ex_flag $\leftarrow$ ex_flag \| (VE \& vxisi_flag)
ex_flag $\leftarrow$ ex_flag | (OE \& ox_flag)
ex_flag $\leftarrow$ ex_flag | (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
end
if( ex_flag =0) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3 , do the following.
For xvmsubasp, do the following.

- Let src1 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X A]$.
- Let src2 be the single-precision floating-point operand in word element i of VSR[XT].
- Let src3 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.

For $\boldsymbol{x v m s u b m s p}$, do the following.

- Let src1 be the single-precision floating-point operand in word element i of VSR[XA].
- Let src2 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.
- Let src3 be the single-precision floating-point operand in word element i of VSR[XT].
src1 is multiplied ${ }^{[1]}$ by src3, producing a product having unbounded range and precision.

See part 1 of Table 127.
src2 is negated and added ${ }^{[2]}$ to the product, producing a sum having unbounded range and precision.

The sum is normalized ${ }^{[3]}$.
See part 2 of Table 127.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into word element $i$ of VSR[XT] in single-precision format.

See Table 106, "Vector Floating-Point Final Result," on page 663.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

FX OX UX XX VXSNAN VXISI VXIMZ

[^5]
## 722

| VSR Data Layout for xvmsub(alm)spsrc1 = VSR[XA] |  |  |  |
| :---: | :---: | :---: | :---: |
| SP | SP | SP | SP |
| src2 = xvmsubasp ? VSR [XT] : VSR[XB] |  |  |  |
| SP | SP | SP | SP |
| src3 = xvmsubasp ? VSR[XB] : VSR[XT] |  |  |  |
| SP | SP | SP | SP |
| tgt $=\mathrm{VSR}[\mathrm{XT}]$ |  |  |  |
| SP | SP | SP | SP |
| 0 |  |  |  |



| src3 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow-$ Infinity | $p \leftarrow-$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $p \leftarrow+$ Infinity | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow+$ Zero | $\mathrm{p} \leftarrow-$ Zero | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow$ src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow+$ Zero | $p \leftarrow+$ Zero | $\mathrm{p} \leftarrow-$ Zero | $p \leftarrow-$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\mathrm{p} \leftarrow-$ Zero | $p \leftarrow-$ Zero | $\mathrm{p} \leftarrow+$ Zero | $p \leftarrow+$ Zero | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow$ Src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $\mathrm{p} \leftarrow-$ Zero | $\mathrm{p} \leftarrow+$ Zero | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow$ src3 | $\begin{aligned} & \mathrm{p} \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{srcl}$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $\begin{aligned} & p \leftarrow \operatorname{src1} \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & P \leftarrow Q(\text { src1 } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & \mathrm{p} \leftarrow Q(\text { src1 } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |


| Part 2: |
| :---: |
| Subtract |
| -Infinity <br> -NZF <br> -Zero <br> +Zero <br> +NZF <br> +Infinity <br>  <br> src1 is a NaN <br>  <br> src1 not a NaN |


| src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\begin{aligned} & \hline \mathrm{V} \leftarrow \mathrm{~d} Q \mathrm{NaN} \\ & \text { vxisi_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\mathrm{V} \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $\mathrm{V} \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \hline v \leftarrow Q(\text { src2 }) \\ & v x \text { snan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $V \leftarrow+$ Infinity | $v \leftarrow S(p, s r c 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow S\left(p, s r^{\prime} 2\right)$ | $\checkmark \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{array}{\|l} \hline \mathrm{V} \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $V \leftarrow+$ Infinity | $\mathrm{v} \leftarrow-\mathrm{SrC2}$ | $v \leftarrow$-Zero | $v \leftarrow$ Rezd | $\mathrm{v} \leftarrow-\mathrm{SrC2}$ | $\checkmark \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l\|} \hline \mathrm{V} \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $V \leftarrow+$ Infinity | $\mathrm{v} \leftarrow-\mathrm{SrC2}$ | $v \leftarrow$ Rezd | $v \leftarrow+$ Zero | $\mathrm{V} \leftarrow-\mathrm{src} 2$ | $V \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \hline \mathrm{V} \leftarrow Q(\text { (src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $V \leftarrow+$ Infinity | $v \leftarrow S(p, s r c 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{v} \leftarrow \mathrm{S}(\mathrm{p}, \mathrm{sc} 22)$ | $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{aligned} & \mathrm{V} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $v \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $\begin{aligned} & \hline \mathrm{v} \leftarrow \mathrm{~d} Q \mathrm{NaN} \\ & \text { vxisi_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{aligned} & v \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\begin{aligned} & v \leftarrow p \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & v \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |

## Explanation:

| src1 | The single-precision floating-point value in word element i of VSR[XA] (where i $\in\{0,1,2,3\}$ ). |
| :---: | :---: |
| src2 | For $\boldsymbol{x v m s u b a s p}$, the single-precision floating-point value in word element $i$ of VSR[XT] (where $i \in\{0,1,2,3\}$ ). For $\boldsymbol{x v m s u b m s p}$, the single-precision floating-point value in word element $i$ of VSR[XB] (where $i \in\{0,1,2,3\}$ ). |
| src3 | For $\boldsymbol{x v m s u b a s p}$, the single-precision floating-point value in word element $i$ of VSR[XB] (where $i \in\{0,1,2,3\}$ ). For $\boldsymbol{x v m s u b m s p}$, the single-precision floating-point value in word element $i$ of VSR[XT] (where $i \in\{0,1,2,3\}$ ). |
| dQNaN | Default quiet NaN (0x7FC0_0000). |
| NZF | Nonzero finite number. |
| Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| $\mathrm{S}(\mathrm{x}, \mathrm{y})$ | Return the normalized sum of floating-point value $x$ and negated floating-point value $y$, having unbounded range and precision. Note: If $x=y, v$ is considered to be an exact-zero-difference result (Rezd). |
| M ( $\mathrm{x}, \mathrm{y}$ ) | Return the normalized product of floating-point value $x$ and floating-point value y , having unbounded range and precision. |
| p | The intermediate product having unbounded range and precision. |
| $v$ | The intermediate result having unbounded range and precision. |

Table 127.Actions for xvmsub(alm)sp

## VSX Vector Multiply Double-Precision XX3-form



Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let src1 be the double-precision floating-point operand in doubleword element $i$ of VSR[XA].

Let src2 be the double-precision floating-point operand in doubleword element $i$ of VSR[XB].
src1 is multiplied ${ }^{[1]}$ by src2, producing a product having unbounded range and precision.

The product is normalized ${ }^{[2]}$.
See Table 128.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into doubleword element $i$ of VSR[XT] in double-precision format.

[^6]|  | src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| -Infinity | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leftarrow+$ Infinity | $\begin{aligned} & \hline v \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{v} \leftarrow \mathrm{dQNaN}$ vximz_flag $\leftarrow 1$ | $v \leftarrow-$ Infinity | $\mathrm{v} \leftarrow-$ Infinity | $\mathrm{v} \leqslant$ src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\mathrm{src} 2) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |
| -NZF | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leftarrow \mathrm{M}(\mathrm{srct} 1, \mathrm{src} 2)$ | $\mathrm{v} \leftarrow+$ Zero | $\mathrm{v} \leftarrow$-Zero | $v \leftarrow M(s \mathrm{sc} 1, \mathrm{src} 2)$ | $v \leftarrow+$ lnfinity | v ¢ STC2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 } 2) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |
| -Zero | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $v \leftarrow+$ Zero | $v \leftarrow+$ Zero | $v \leftarrow-$ Zero | $v \leftarrow$-Zero | $\mathrm{V} \leftarrow \mathrm{dQNaN}$ vximz_flag $\leftarrow 1$ | v ¢ STC2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srco2) } \\ & \mathrm{vxsnan} f l a g \leftarrow 1 \end{aligned}$ |
| $\overline{\text { ] }}$ +Zero | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $v \leftarrow-$ Zero | $\mathrm{v} \leftarrow$-Zero | $v \leftarrow+$ Zero | $v \leftarrow+$ Zero | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vximz_lag } \leftarrow 1 \end{aligned}$ | v ¢ SrC2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srco2) } \\ & \mathrm{vxsnan} \mathrm{flag} \leftarrow 1 \end{aligned}$ |
| +NZF | $\mathrm{v} \leftarrow-$ Infinity | $\mathrm{v} \leftarrow \mathrm{M}$ (srct 1 ,sc2) | $\mathrm{v} \leftarrow-$ Zero | $v \leftarrow+$ Zero | $v \leftarrow M(s \mathrm{sc} 1, \mathrm{src} 2)$ | $v \leftarrow+$ Infinity | v ¢ SrC2 | $\begin{aligned} & \hline \begin{array}{l} v \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \end{array} \end{aligned}$ |
| +Infinity | $\mathrm{v} \leftarrow-$ Infinity | $\mathrm{v} \leftarrow+$ Infinity | $\begin{aligned} & \hline v \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{v} \leftarrow \mathrm{dQNaN}$ vximz_flag $\leftarrow 1$ | $v \leftarrow+$ lnfinity | $v \leftarrow+$ lnfinity | v ¢ Src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\mathrm{src} 2) \\ & \text { vxsnan_flag } \leqslant 1 \end{aligned}$ |
| QNaN | $\mathrm{v} \leqslant$ Srct | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant$ Src1 | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant$ SrC1 | $\begin{aligned} & \mathrm{v} \leftarrow \text { Src1 } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| SNaN | $\begin{aligned} & \mathrm{v} \leftarrow Q(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline \begin{array}{l} \mathrm{v} \leftarrow Q(\text { srcil }) \\ \text { vxsnan_flag } \leftarrow 1 \end{array} \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{V} \leftarrow Q(\text { srcic }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srcil }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline \begin{array}{l} v \leftarrow Q(\text { srcil }) \\ v x s n a n \_f l a g \leftarrow 1 \end{array} \end{aligned}$ | $\begin{aligned} & \mathrm{V} \leftarrow Q(\text { srcic }) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srcil }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |


| Explanation: |  |
| :---: | :---: |
| src1 | The double-precision floating-point value in doubleword element i of VSR[XA] (where $\mathrm{i} \in\{0,1\}$ ). |
| src2 | The double-precision floating-point value in doubleword element i of VSR[XB] (where $\mathrm{i} \in\{0,1\}$ ). |
| dQNaN | Default quiet NaN (0x7FF8_0000_0000_0000). |
| NZF | Nonzero finite number. |
| $\mathrm{M}(\mathrm{x}, \mathrm{y})$ | Return the normalized product of floating-point value $x$ and floating-point value y , having unbounded range and precision. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| $v$ | The intermediate result having unbounded signficand precision and unbounded exponent range. |

Table 128.Actions for xvmuldp

## VSX Vector Multiply Single-Precision XX3-form


i=0 to 127 by 32
reset_xflags()
$\operatorname{src1} \quad \leftarrow \operatorname{VSR}[X A]\{i: i+31\}$
$\operatorname{src3} \quad \leftarrow \operatorname{VSR}[X B]\{i: i+31\}$
v\{0:inf\} $\leftarrow$ MultiplySP(src1,src3)
result $\{\mathrm{i}: i+31\} \leftarrow$ RoundToSP $(R N, v)$
if(vxsnan_flag) then SetFX(VXSNAN)
if(vximz_flag) then SetFX(VXIMZ)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(xx_flag) then SetFX(XX)
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vximz_flag)
ex_flag $\leftarrow$ ex_flag | (OE \& ox_flag)
ex_flag $\leftarrow$ ex_flag | (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
if( ex_flag = 0 ) then VSR $[\mathrm{XT}] \leftarrow$ result

Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3, do the following. Let src1 be the single-precision floating-point operand in word element i of VSR[XA].

Let src2 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.
src1 is multiplied ${ }^{[1]}$ by src2, producing a product having unbounded range and precision.

The product is normalized ${ }^{[2]}$.
See Table 129.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into word element i of $\mathrm{VSR}[\mathrm{XT}]$ in single-precision format.

[^7]|  | src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| -Infinity | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leftarrow+$ Infinity | $\begin{aligned} & \hline v \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{v} \leftarrow \mathrm{dQNaN}$ vximz_flag $\leftarrow 1$ | $v \leftarrow-$ Infinity | $\mathrm{v} \leftarrow-$ Infinity | $\mathrm{v} \leqslant$ src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\mathrm{src} 2) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |
| -NZF | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leftarrow \mathrm{M}(\mathrm{srct} 1, \mathrm{src} 2)$ | $\mathrm{v} \leftarrow+$ Zero | $\mathrm{v} \leftarrow$-Zero | $v \leftarrow M(s \mathrm{sc} 1, \mathrm{src} 2)$ | $v \leftarrow+$ lnfinity | v ¢ STC2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 } 2) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |
| -Zero | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $v \leftarrow+$ Zero | $v \leftarrow+$ Zero | $v \leftarrow-$ Zero | $v \leftarrow$-Zero | $\mathrm{V} \leftarrow \mathrm{dQNaN}$ vximz_flag $\leftarrow 1$ | v ¢ STC2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srco2) } \\ & \mathrm{vxsnan} f l a g \leftarrow 1 \end{aligned}$ |
| $\overline{\text { ] }}$ +Zero | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $v \leftarrow-$ Zero | $\mathrm{v} \leftarrow$-Zero | $v \leftarrow+$ Zero | $v \leftarrow+$ Zero | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vximz_lag } \leftarrow 1 \end{aligned}$ | v ¢ SrC2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srco2) } \\ & \mathrm{vxsnan} \mathrm{flag} \leftarrow 1 \end{aligned}$ |
| +NZF | $\mathrm{v} \leftarrow-$ Infinity | $\mathrm{v} \leftarrow \mathrm{M}$ (srct 1 ,sc2) | $\mathrm{v} \leftarrow-$ Zero | $v \leftarrow+$ Zero | $v \leftarrow M(s \mathrm{sc} 1, \mathrm{src} 2)$ | $v \leftarrow+$ Infinity | v ¢ SrC2 | $\begin{aligned} & \hline \begin{array}{l} v \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \end{array} \end{aligned}$ |
| +Infinity | $\mathrm{v} \leftarrow-$ Infinity | $\mathrm{v} \leftarrow+$ Infinity | $\begin{aligned} & \hline v \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{v} \leftarrow \mathrm{dQNaN}$ vximz_flag $\leftarrow 1$ | $v \leftarrow+$ lnfinity | $v \leftarrow+$ lnfinity | v ¢ Src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\mathrm{src} 2) \\ & \text { vxsnan_flag } \leqslant 1 \end{aligned}$ |
| QNaN | $\mathrm{v} \leqslant$ Srct | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant$ Src1 | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant$ SrC1 | $\begin{aligned} & \mathrm{v} \leftarrow \text { Src1 } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| SNaN | $\begin{aligned} & \mathrm{v} \leftarrow Q(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline \begin{array}{l} \mathrm{v} \leftarrow Q(\text { srcil }) \\ \text { vxsnan_flag } \leftarrow 1 \end{array} \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{V} \leftarrow Q(\text { srcic }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srcil }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline \begin{array}{l} v \leftarrow Q(\text { srcil }) \\ v x s n a n \_f l a g \leftarrow 1 \end{array} \end{aligned}$ | $\begin{aligned} & \mathrm{V} \leftarrow Q(\text { srcic }) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srcil }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |


| Explanation: |  |
| :--- | :--- | :--- |
| $\operatorname{src} 1$ | The single-precision floating-point value in word element i of $\operatorname{VSR}[\mathrm{XA}]$ (where $\mathrm{i} \in\{0,1,2,3\})$. |
| src 2 | The single-precision floating-point value in word element i of $\operatorname{VSR}[\mathrm{XB}]$ (where $\mathrm{i} \in\{0,1,2,3\})$. |
| dQNaN | Default quiet $\mathrm{NaN}\left(0 \times 7 \mathrm{FC} 0 \_0000\right)$. |
| NZF | Nonzero finite number. |
| $\mathrm{M}(\mathrm{x}, \mathrm{y})$ | Return the normalized product of floating-point value x and floating-point value y, , having unbounded range and precision. |
| $\mathrm{Q}(\mathrm{x})$ | Return a QNaN with the payload of x. |
| v | The intermediate result having unbounded signficand precision and unbounded exponent range. |

Table 129.Actions for xvmulsp

## VSX Vector Negative Absolute Double-Precision XX2-form

| xvnabs |  | XT, XB |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ${ }_{0} 60$ | ${ }^{6}$ | ${ }_{11} \quad \text { III }$ | 16 | B | 21 | 489 | $\left\|\begin{array}{l}\text { BX } \\ 30 \times \\ 30\end{array}\right\|$ |

```
XT}\leftarrowTX|
XB}\leftarrowBX|
do i=0 to 127 by 64
    VSR[XT]{i:i+63} < 0b1 || VSR[XB]{i+1:i+63}
end
```

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. The contents of doubleword element $i$ of VSR[XB], with bit 0 set to 1 , is placed into doubleword element i of VSR[XT].

Special Registers Altered
None

VSR Data Layout for xvnabsdp
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| DP | DP |
| :---: | :---: |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

|  | DP |
| :--- | :--- |
| 0 | 64 |

## VSX Vector Negative Absolute

 Single-Precision XX2-formxvnabssp $\quad \mathrm{XT}, \mathrm{XB}$

| ${ }_{0} 60$ | 6 | T | 11 | III | 16 | B | 21 | 425 | BXXTX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |



```
XB}\leftarrowBX|
do i=0 to 127 by 32
    VSR[XT]{i:i+31} < 0b1 || VSR[XB]{i+1:i+31}
end
```

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3, do the following. The contents of word element $i$ of $\operatorname{VSR}[X B]$, with bit 0 set to 1 , is placed into word element $i$ of VSR[XT].

Special Registers Altered None

VSR Data Layout for xvnabssp
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| $S P$ | $S P$ | $S P$ | $S P$ |
| :--- | :--- | :--- | :--- |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

| SP | SP |  | SP |  |
| :--- | :--- | :--- | :--- | :--- |
| 0 | 32 | SP |  |  |
| 0 |  |  |  |  |

## Version 3.0

VSX Vector Negate Double-Precision XX2-form
Xvnegdp

| 60 |  | T |  | IIII |  | B |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 60 |  |  | 505 | $3 X T X$ |  |  |
| 3031 |  |  |  |  |  |  |

```
XT}\leftarrowTX||
XB}\leftarrowBX|
do i=0 to 127 by 64
    VSR[XT]{i:i+63}}\leftarrow~VSR[XB]{i} || VSR[XB]{i+1:i+63
end
```

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. The contents of doubleword element $i$ of VSR[XB], with bit 0 complemented, is placed into doubleword element i of VSR[XT].

## Special Registers Altered None

| VSR Data Layout for xvnegdp |
| :--- |
| src $=\mathrm{VSR}[\mathrm{XB}]$ |
| DP DP |

tgt $=\operatorname{VSR}[\mathrm{XT}]$

|  | DP |
| :--- | :--- |
| 0 | 64 |

VSX Vector Negate Single-Precision XX2-form

$X T \leftarrow T X \| T$
$X B \leftarrow B X \| B$

```
do i=0 to 127 by 32
    VSR[XT]{i:i+31}}\leftarrow~\mathrm{ VSR[XB]{i} || VSR[XB]{i+1:i+31}
end
```

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3, do the following. The contents of word element $i$ of $\operatorname{VSR}[X B]$, with bit 0 complemented, is placed into word element $i$ of VSR[XT].

## Special Registers Altered

None

## VSR Data Layout for xvnegsp

src = VSR[XB]

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

| SP | SP |  | SP | SP |  |
| :--- | :--- | :--- | :--- | :--- | :---: |
| 0 | 32 | 64 | 96 | 127 |  |

## VSX Vector Negative Multiply-Add Double-Precision XX3-form

xvnmaddadp $\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}$

| ${ }_{0} 60$ | 6 |  | A | 16 | B | 21 | 225 | $\left\lvert\, \begin{aligned} & \text { AXX } \\ & 29 \times 1 \text { TX } \\ & 29031\end{aligned}\right.$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| xvnmaddmdp |  | XT, XA, XB |  |  |  |  |  |  |


| 60 | 6 | T | 11 | A | 16 | B | 21 | 233 | $\left[\left.\begin{array}{l} \operatorname{Ax} \times 8 \times 1 \mathrm{TX} \\ 293031 \end{array} \right\rvert\,\right.$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |


| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0$ 0b0 |

do $i=0$ to 127 by 64
reset_xflags()
$\operatorname{src} 1 \leftarrow \operatorname{VSR}[X A]\{i: i+63\}$
src2 $\leftarrow$ "xvnmaddadp" ? VSR[XT]\{i:i+63\}: VSR[XB]\{i:i+63\}
src3 $\leftarrow$ "xvnmaddadp" ? VSR[XB]\{i:i+63\} : VSR[XT]\{i:i+63\}
v\{0:inf\} $\leftarrow$ MultiplyAddDP(src1,src3,src2)
result $\{\mathrm{i}: i+63\} \leftarrow \operatorname{NegateDP}($ RoundToDP(RN, v))
if(vxsnan_flag) then SetFX(VXSNAN)
if(vximz_flag) then SetFX(VXIMZ)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(xx_flag) then $\operatorname{SetFX}(X X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vximz_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vxisi_flag)
ex_flag $\leftarrow$ ex_flag | (OE \& ox_flag)
ex_flag $\leftarrow$ ex_flag | (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
end
if( ex_flag $=0$ ) then VSR $[X T] \leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following.
For xvnmaddadp, do the following.

- Let src1 be the double-precision floating-point operand in doubleword element i of VSR[XA].
- Let src2 be the double-precision floating-point operand in doubleword element i of VSR[XT].
- Let src3 be the double-precision floating-point operand in doubleword element i of VSR[XB].

For xvnmaddmdp, do the following.

- Let src1 be the double-precision floating-point operand in doubleword element i of VSR[XA].
- Let src2 be the double-precision floating-point operand in doubleword element i of VSR[XB].
- Let src3 be the double-precision floating-point operand in doubleword element i of VSR[XT].
src1 is multiplied ${ }^{[1]}$ by src3, producing a product having unbounded range and precision.

See part 1 of Table 130.
src2 is added ${ }^{[2]}$ to the product, producing a sum having unbounded range and precision.

The sum is normalized ${ }^{[3]}$.
See part 2 of Table 130.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is negated and placed into doubleword element $i$ of $\operatorname{VSR}[\mathrm{XT}]$ in double-precision format.

See Table 131, "Vector Floating-Point Final Result with Negation," on page 734.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

Special Registers Altered<br>FX OX UX XX VXSNAN VXISI VXIMZ

## Version 3.0




| src3 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow-$ Infinity | $p \leftarrow-$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow+$ Infinity | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow+$ Zero | $p \leftarrow+$ Zero | $p \leftarrow-$ Zero | $p \leftarrow-$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow$ Src3 | $\begin{aligned} & p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow-$ Zero | $p \leftarrow-$ Zero | $p \leftarrow+$ Zero | $p \leftarrow+$ Zero | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow+$ Infinity | $\begin{array}{\|l\|} \hline p \leftarrow d Q N a N \\ \text { vximz_flag } \leftarrow 1 \\ \hline \end{array}$ | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{array}{\|l} \hline p \leftarrow Q(\text { src } 3) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{srcl}$ | $p \leftarrow \operatorname{src} 1$ | $\begin{aligned} & p \leftarrow \operatorname{src} 1 \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & p \leftarrow Q(\text { src } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & P \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |



| src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $V \leftarrow-$ Infinity | $\mathrm{V} \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $\mathrm{V} \leftarrow-$ Infinity | $\begin{array}{\|l} \hline \mathrm{v} \leftarrow \mathrm{~d} Q \mathrm{NaN} \\ \text { vxisi_flag } \leftarrow 1 \end{array}$ | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{array}{\|l\|} \hline \mathrm{v} \leftarrow Q \text { (src2) } \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{A}(\mathrm{p}, \mathrm{src} 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow A(p, s r c 2)$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l\|} \hline v \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\mathrm{V} \leftarrow$-Zero | $v \leftarrow$ Rezd | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{array}{\|l\|} \hline v \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $\mathrm{V} \leftarrow$ - Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $v \leftarrow$ Rezd | $\mathrm{V} \leftarrow+$ Zero | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{V} \leftarrow Q \text { (src2) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{A}(\mathrm{p}, \mathrm{src} 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow A(p, s r c 2)$ | $\mathrm{V} \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & v x \text { vnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\begin{array}{\|l\|} \hline \mathrm{V} \leftarrow \mathrm{~d} Q \mathrm{NaN} \\ \text { vxisi_flag } \leftarrow 1 \end{array}$ | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{aligned} & \hline v \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\begin{aligned} & v \leftarrow p \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \hline v \leftarrow Q(\text { src2 }) \\ & v x \text { snan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |


| Explanation: |  |
| :---: | :---: |
| src1 | The double-precision floating-point value in doubleword element $i$ of VSR[XA] (where $i \in\{0,1\}$ ). |
| src2 | For $\boldsymbol{x v n m a d d a d p}$, the double-precision floating-point value in doubleword element $i$ of VSR[XT] (where $i \in\{0,1\}$ ). For $\boldsymbol{x v n m a d d m d p}$, the double-precision floating-point value in doubleword element $i$ of VSR[XB] (where $i \in\{0,1\}$ ). |
| src3 | For $\boldsymbol{x v n m a d d a d p ,}$ the double-precision floating-point value in doubleword element $i$ of VSR[XB] (where $i \in\{0,1\}$ ). For $\boldsymbol{x v n m a d d m d p}$, the double-precision floating-point value in doubleword element $i$ of VSR[XT] (where $i \in\{0,1\}$ ). |
| dQNaN | Default quiet NaN (0x7FF8_0000_0000_0000). |
| NZF | Nonzero finite number. |
| Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| A(x,y) | Return the normalized sum of floating-point value $x$ and floating-point value $y$, having unbounded range and precision. Note: If $x=-y$, $v$ is considered to be an exact-zero-difference result (Rezd). |
| M(x,y) | Return the product of floating-point value $x$ and floating-point value y , having unbounded range and precision. |
| p | The intermediate product having unbounded range and precision. |
| $v$ | The intermediate result having unbounded range and precision. |

Table 130.Actions for xvnmadd(alm)dp

## Version 3.0



Table 131.Vector Floating-Point Final Result with Negation


Table 131.Vector Floating-Point Final Result with Negation (Continued)

## VSX Vector Negative Multiply-Add Single-Precision XX3-form

$$
\text { xvnmaddasp } \quad X T, X A, X B
$$


xvnmaddmsp $\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}$

| 60 |  | T |  | A |  | B |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |


| $X T$ | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

do $\mathrm{i}=0$ to 127 by 32
reset_xflags()
src1 $\leftarrow$ VSR[XA]\{i:i+31\}
src2 $\leftarrow$ "xvnmaddasp" ? VSR $[X T]\{i: i+31\}: V S R[X B]\{i: i+31\}$
src3 $\leftarrow$ "xvnmaddasp" ? VSR[XB]\{i:i+31\}:VSR[XT]\{i:i+31\}
v\{0:inf\} $\leftarrow$ MultiplyAddSP(src1,src3,src2)
result $\{i: i+31\} \leftarrow \operatorname{NegateSP}($ RoundToSP(RN,v))
if(vxsnan_flag) then SetFX(vXSNAN)
if(vximz_flag) then SetFX(vXIMZ)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then $\operatorname{SetFX}(0 X)$
if(ux_flag) then SetFX(UX)
if(xx_flag) then $\operatorname{SetFX}(x X)$
ex_flag $\leftarrow$ ex_flag \| (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag \| (VE \& vximz_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vxisi_flag)
ex_flag $\leftarrow$ ex_flag | ( $O E \&$ ox_flag)
ex_flag $\leftarrow$ ex_flag | (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
end

$$
\text { if( ex_flag = } 0 \text { ) then VSR }[X T] \leftarrow \text { result }
$$

Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3 , do the following.
For xvnmaddasp, do the following.

- Let src1 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X A]$.
- Let src2 be the single-precision floating-point operand in word element $i$ of VSR[XT].
- Let src3 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.

For xvnmaddmsp, do the following.

- Let src1 be the single-precision floating-point operand in word element i of VSR[XA].
- Let src2 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.
- Let src3 be the single-precision floating-point operand in word element i of VSR[XT].
src1 is multiplied ${ }^{[1]}$ by src3, producing a product having unbounded range and precision.

See part 1 of Table 132.
src2 is added ${ }^{[2]}$ to the product, producing a sum having unbounded range and precision.

The sum is normalized ${ }^{[3]}$.
See part 2 of Table 132.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is negated and placed into word element $i$ of VSR[XT] in single-precision format.

See Table 131, "Vector Floating-Point Final Result with Negation," on page 734.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

FX OX UX XX VXSNAN VXISI VXIMZ

[^8]


| src3 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & \mathrm{p} \leftarrow \mathrm{dQNaN} \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow-$ Infinity | $p \leftarrow-$ Infinity | $p \leftarrow$ Src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $p \leftarrow+$ Infinity | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $\mathrm{p} \leftarrow \mathrm{SrCl}$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & p \leftarrow Q(\operatorname{src} 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{p} \leftarrow+$ Zero | $\mathrm{p} \leftarrow+$ Zero | $\mathrm{p} \leftarrow-$ Zero | $p \leftarrow-$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\mathrm{p} \leftarrow$-Zero | $\mathrm{p} \leftarrow$-Zero | $\mathrm{p} \leftarrow+$ Zero | $\mathrm{p} \leftarrow+$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & p \leftarrow Q \text { (src3) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $\mathrm{p} \leftarrow \mathrm{SrCl}$ | $\mathrm{p} \leftarrow \mathrm{src} 1$ | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow$ Src3 | $\begin{aligned} & p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & p \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $\mathrm{p} \leftarrow \mathrm{SrC1}$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $\begin{aligned} & p \leftarrow \operatorname{src1} \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & \mathrm{p} \leftarrow Q(\text { src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q \text { (src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q \text { (src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q \text { (src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |



| src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\mathrm{V} \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $\begin{array}{\|l\|} \hline v \leftarrow d Q N a N \\ \text { vxisi_flag } \leftarrow 1 \end{array}$ | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \hline v \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $v \leftarrow A(p, s r c 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow A\left(p, s c^{\prime} 2\right)$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{v} \leftarrow \mathrm{SrC2}$ | $\mathrm{V} \leftarrow$-Zero | $v \leftarrow$ Rezd | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $\begin{array}{\|l\|} \hline v \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $v \leftarrow$ Rezd | $\mathrm{v} \leftarrow+$ Zero | $\mathrm{V} \leftarrow \mathrm{src} 2$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\mathrm{V} \leftarrow-$ Infinity | $v \leftarrow A(p, s r c 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow A(p, s r c 2)$ | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \hline v \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{dQNaN} \\ & \text { vxisi_flag } \leftarrow 1 \end{aligned}$ | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\begin{aligned} & v \leftarrow p \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |

## Explanation:

| src1 | The single-precision floating-point value in word element i of VSR[XA] (where $i \in\{0,1,2,3\}$ ). |
| :---: | :---: |
| src2 | For xvnmaddasp, the single-precision floating-point value in word element $i$ of VSR[XT] (where $i \in\{0,1,2,3\}$ ). For $\boldsymbol{x v n m a d d m s p}$, the single-precision floating-point value in word element $i$ of VSR[XB] (where $i \in\{0,1,2,3\}$ ). |
| src3 | For xvnmaddasp, the single-precision floating-point value in word element i of VSR[XB] (where $i \in\{0,1,2,3\}$ ). For $\boldsymbol{x v n m a d d m s p , ~ t h e ~ s i n g l e - p r e c i s i o n ~ f l o a t i n g - p o i n t ~ v a l u e ~ i n ~ w o r d ~ e l e m e n t ~} i$ of VSR[XT] (where $i \in\{0,1,2,3\}$ ). |
| dQNaN | Default quiet NaN (0x7FC0_0000). |
| NZF | Nonzero finite number. |
| Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| A(x,y) | Return the normalized sum of floating-point value $x$ and floating-point value $y$, having unbounded range and precision. Note: If $x=-y, v$ is considered to be an exact-zero-difference result (Rezd). |
| M ( $\mathrm{x}, \mathrm{y}$ ) | Return the normalized product of floating-point value $x$ and floating-point value y , having unbounded range and precision. |
| p | The intermediate product having unbounded range and precision. |
| v | The intermediate result having unbounded range and precision. |

Table 132.Actions for xvnmadd(alm)sp

## VSX Vector Negative Multiply-Subtract Double-Precision XX3-form

xvnmsubadp $\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}$

| 60 | 6 |  | A |  | B | 21 | 241 | $\left\lvert\, \begin{aligned} & \text { Ax } \\ & 29 \times 1 \text { PV } \\ & 2031\end{aligned}\right.$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| xvnmsubmdp |  |  |  |  |  |  |  |  |



| XT | $\leftarrow$ TX \\| $\\|$ T |
| :--- | :--- |
| XA | $\leftarrow$ AX \\| A |
| XB | $\leftarrow$ BX \\| |

if( ex_flag = 0) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following.
For xvmsubadp, do the following.

- Let src1 be the double-precision floating-point operand in doubleword element i of VSR[XA].
- Let src2 be the double-precision floating-point operand in doubleword element i of VSR[XT].
- Let src3 be the double-precision floating-point operand in doubleword element i of VSR[XB].

For xvmsubmdp, do the following.

- Let src1 be the double-precision floating-point operand in doubleword element i of VSR[XA].
- Let src2 be the double-precision floating-point operand in doubleword element i of VSR[XB].
- Let src3 be the double-precision floating-point operand in doubleword element i of VSR[XT].
src1 is multiplied ${ }^{[1]}$ by src3, producing a product having unbounded range and precision.

See part 1 of Table 133.
src2 is negated and added ${ }^{[2]}$ to the product, producing a sum having unbounded range and precision.

The sum is normalized ${ }^{[3]}$.
See part 2 of Table 133.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is negated and placed into doubleword element $i$ of VSR[XT] in double-precision format.

See Table 131, "Vector Floating-Point Final Result with Negation," on page 734.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

FX OX UX XX VXSNAN VXISI VXIMZ

## Version 3.0




| src3 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow-$ Infinity | $p \leftarrow-$ Infinity | $p \leftarrow$ src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $p \leftarrow+$ Infinity | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow$ Srcl | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow$ src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow+$ Zero | $\mathrm{p} \leftarrow+$ Zero | $p \leftarrow-$ Zero | $p \leftarrow-$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow$ Src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow-$ Zero | $p \leftarrow$-Zero | $\mathrm{p} \leftarrow+$ Zero | $p \leftarrow+$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow$ src3 | $\begin{aligned} & p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & \hline p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \\ & \hline \end{aligned}$ | $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $p \leftarrow$ src3 | $\begin{array}{\|l\|} \hline p \leftarrow Q(\text { src3 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $p \leftarrow \operatorname{src} 1$ | $\mathrm{p} \leftarrow \mathrm{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow$ Srcl | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{srcl}$ | $\begin{aligned} & p \leftarrow \operatorname{src1} \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow Q(\operatorname{src} 1)$ <br> vxsnan_flag $\leftarrow 1$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & P \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & P \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow Q(\operatorname{src} 1)$ <br> vxsnan_flag $\leftarrow 1$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |



| src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\begin{array}{\|l\|} \hline \begin{array}{l} \mathrm{v} \leftarrow \mathrm{~d} Q N a N \\ \mathrm{vxisif} \text { flag } \leftarrow 1 \end{array} \\ \hline \end{array}$ | $\mathrm{V} \leftarrow-$-nfinity | $\mathrm{V} \leqslant-$ - lninity | $\mathrm{V} \leqslant-$ - lnfinity | $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow-$ Infinity | v ¢ Src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 } 2) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |
| $\mathrm{v} \leftarrow+$ lnfinity | $v \leftarrow S(p, s \mathrm{sc} 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow S(p, s \mathrm{sc} 2)$ | $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leqslant$ Src2 | $\begin{aligned} & \hline \mathrm{V} \leftarrow Q(\text { src } 2) \\ & \mathrm{vxsnn} \text { _flag } \leftarrow 1 \end{aligned}$ |
| $\mathrm{v} \leftarrow+$ lnfinity | $\mathrm{v} \leftarrow-\mathrm{SrC2}$ | $v \leftarrow-$ Zero | $v \leftarrow$ Rezd | $\mathrm{v} \leftarrow-\mathrm{SrC2}$ | $\mathrm{v} \leftarrow-$ Infinity | v \& SrC2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src } 2) \\ & \mathrm{vxsnn} \text { _flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow+$ lnfinity | $\mathrm{v} \leftarrow-\mathrm{SrC2}$ | $v \leftarrow$ Rezd | $v \leftarrow+$ Zero | $\mathrm{v} \leftarrow-\mathrm{SrC2}$ | $\mathrm{V} \leftarrow-$ Infinity | v \& SrC2 | $\begin{aligned} & \hline v \leftarrow Q(\text { src2 }) \\ & v x s n a n \_f l a g \leftarrow 1 \end{aligned}$ |
| $\mathrm{V} \leftarrow+$ lnfinity | $v \leftarrow S(p, s \mathrm{sc} 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow S(p, s \mathrm{sr} 2)$ | $\mathrm{V} \leftarrow-$ Infinity | V \& Src2 | $\begin{aligned} & \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src } 2) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow+$ lnfinity | $v \leftarrow+$ lnfinity | $v \leftarrow+$ lnfinity | $v \leftarrow+$ lnfinity | $\mathrm{V} \leftarrow+$ lnfinity | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{~d} \text { NaN } \\ & \text { vxisi_flag } \leftarrow 1 \end{aligned}$ | v \& SrC2 | $\begin{aligned} & \hline \mathrm{v} \leftarrow Q(\text { srč2) } \\ & \mathrm{vxsnnan} \text { flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ <br> vxsnan_flag $\leftarrow 1$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | V \& SrC2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src2 } 2) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |


| Explanation: |  |
| :---: | :---: |
| src1 | The double-precision floating-point value in doubleword element $i$ of VSR[XA] (where $i \in\{0,1\}$ ). |
| src2 | For $\boldsymbol{x} v n m s u b a d p$, the double-precision floating-point value in doubleword element $i$ of $\operatorname{VSR}[X T]$ (where $i \in\{0,1\}$ ). For $\boldsymbol{x} v n m s u b m d p$, the double-precision floating-point value in doubleword element $i$ of VSR[XB] (where $i \in\{0,1\}$ ). |
| src3 | For $\boldsymbol{x v n m s u b a d p}$, the double-precision floating-point value in doubleword element $i$ of VSR[XB] (where $i \in\{0,1\}$ ). For $\boldsymbol{x v n m s u b m d p}$, the double-precision floating-point value in doubleword element $i$ of VSR[XT] (where $i \in\{0,1\}$ ). |
| dQNaN | Default quiet NaN (0x7FF8_0000_0000_0000). |
| NZF | Nonzero finite number. |
| Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| $S(x, y)$ | Return the normalized sum of floating-point value $x$ and negated floating-point value $y$, having unbounded range and precision. Note: If $x=-y$, $v$ is considered to be an exact-zero-difference result (Rezd). |
| M ( $\mathrm{x}, \mathrm{y}$ ) | Return the normalized product of floating-point value $x$ and floating-point value y , having unbounded range and precision. |
| p | The intermediate product having unbounded range and precision. |
| v | The intermediate result having unbounded range and precision. |

Table 133.Actions for xvnmsub(alm)dp

## VSX Vector Negative Multiply-Subtract Single-Precision XX3-form

$$
\text { xvnmsubasp } \quad \text { XT,XA,XB }
$$

| 60 | 6 | T | 11 | A | 16 | B | 21 | 209 | $\|a x\| B X T x \mid$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

xvnmsubmsp $\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}$


| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

do i=0 to 127 by 32
reset_xflags()
$\operatorname{src} 1 \leftarrow \operatorname{VSR}[X A]\{i: i+31\}$
$\operatorname{src} 2 \leftarrow$ "xvnmsubasp" ? VSR $[X T]\{i: i+31\}: V S R[X B]\{i: i+31\}$
src3 $\leftarrow$ "xvnmsubasp" ? VSR[XB]\{i:i+31\}:VSR[XT]\{i:i+31\}
v\{0:inf\} $\leftarrow$ MultiplyAddSP(src1,src3,NegateSP(src2))
result $\{i: i+31\} \leftarrow \operatorname{NegateSP}($ RoundToSP(RN,v))
if(vxsnan_flag) then SetFX(vXSNAN)
if(vximz_flag) then SetFX(vXIMZ)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then $\operatorname{SetFX}(0 X)$
if(ux_flag) then SetFX(UX)
if(xx_flag) then $\operatorname{SetFX}(x X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag \| (VE \& vximz_flag)
ex_flag $\leftarrow$ ex_flag \| (VE \& vxisi_flag)
ex_flag $\leftarrow$ ex_flag | ( $0 \mathrm{E} \&$ ox_flag)
ex_flag $\leftarrow$ ex_flag | (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
end

$$
\text { if( ex_flag = } 0 \text { ) then VSR }[X T] \leftarrow \text { result }
$$

Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3 , do the following.
For xvnmsubasp, do the following.

- Let src1 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X A]$.
- Let src2 be the single-precision floating-point operand in word element i of VSR[XT].
- Let src3 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.

For xvnmsubmsp, do the following.

- Let src1 be the single-precision floating-point operand in word element i of VSR[XA].
- Let src2 be the single-precision floating-point operand in word element $i$ of $\operatorname{VSR}[X B]$.
- Let src3 be the single-precision floating-point operand in word element i of VSR[XT].
src1 is multiplied ${ }^{[1]}$ by src3, producing a product having unbounded range and precision.

See part 1 of Table 134.
src2 is negated and added ${ }^{[2]}$ to the product, producing a sum having unbounded range and precision.

The sum is normalized ${ }^{[3]}$.
See part 2 of Table 134.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is negated and placed into word element $i$ of VSR[XT] in single-precision format.

See Table 131, "Vector Floating-Point Final Result with Negation," on page 734.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

FX OX UX XX VXSNAN VXISI VXIMZ

[^9]
## 742




| src3 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow-$ Infinity | $p \leftarrow-$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow+$ Infinity | $p \leftarrow M(s r c 1, s r c 3)$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow$ src3 | $\begin{aligned} & \hline p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \\ & \hline \end{aligned}$ |
| $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{p} \leftarrow+$ Zero | $\mathrm{p} \leftarrow+$ Zero | $p \leftarrow-$ Zero | $\mathrm{p} \leftarrow$-Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \mathrm{p} \leftarrow Q(\text { src } 3) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow-$ Zero | $\mathrm{p} \leftarrow$-Zero | $\mathrm{p} \leftarrow+$ Zero | $\mathrm{p} \leftarrow+$ Zero | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & p \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow M(\operatorname{src} 1, \operatorname{src} 3)$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow M(\operatorname{src} 1, \mathrm{src} 3)$ | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \mathrm{p} \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow-$ Infinity | $p \leftarrow+$ Infinity | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow d Q N a N \\ & \text { vximz_flag } \leftarrow 1 \end{aligned}$ | $p \leftarrow+$ Infinity | $p \leftarrow+$ Infinity | $p \leftarrow \operatorname{src} 3$ | $\begin{aligned} & \mathrm{p} \leftarrow Q(\text { src3 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $p \leftarrow \operatorname{src} 1$ | $\begin{aligned} & p \leftarrow \operatorname{src1} \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{P} \leftarrow \mathrm{Q}(\text { src } 1) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline p \leftarrow Q(\text { src1) } \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & p \leftarrow Q(\text { src1 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |



| src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\mathrm{v} \leftarrow \mathrm{d}$ QNaN vxisi_flag $\leftarrow 1$ | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $V \leftarrow-$ Infinity | $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l} \hline \mathrm{V} \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $\mathrm{V} \leftarrow+$ Infinity | $v \leftarrow S(p, s r c 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow S(p, s r c 2)$ | $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l} \hline \mathrm{v} \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $V \leftarrow+$ Infinity | $\mathrm{v} \leftarrow-\mathrm{src} 2$ | $\mathrm{v} \leftarrow$-Zero | $v \leftarrow$ Rezd | $\mathrm{V} \leftarrow-\mathrm{src} 2$ | $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l\|} \hline \mathrm{V} \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $V \leftarrow+$ Infinity | $\mathrm{v} \leftarrow-\mathrm{src} 2$ | $v \leftarrow$ Rezd | $v \leftarrow+$ Zero | $\mathrm{V} \leftarrow-\mathrm{src} 2$ | $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l} \hline \mathrm{V} \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{S}(\mathrm{p}, \mathrm{src} 2)$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{V} \leftarrow \mathrm{S}(\mathrm{p}, \mathrm{src} 2)$ | $\mathrm{V} \leftarrow-$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC2}$ | $\begin{array}{\|l} \hline \mathrm{V} \leftarrow Q(\text { src2 }) \\ \text { vxsnan_flag } \leftarrow 1 \\ \hline \end{array}$ |
| $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $v \leftarrow+$ Infinity | $V \leftarrow+$ Infinity | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vxisi_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{v} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & \mathrm{V} \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\begin{aligned} & v \leftarrow p \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |
| $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $v \leftarrow p$ | $\mathrm{v} \leftarrow \mathrm{SrC2}$ | $\begin{aligned} & v \leftarrow Q(\text { src2 }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ |

## Explanation:

| src1 | The single-precision floating-point value in word element i of VSR[XA] (where i $\in\{0,1,2,3\}$ ). |
| :---: | :---: |
| src2 | The single-precision floating-point value in word element i of VSR[XT] (where $i \in\{0,1,2,3\}$ ). |
| src3 | The single-precision floating-point value in word element i of VSR[XB] (where $i \in\{0,1,2,3\}$ ). |
| dQNaN | Default quiet NaN (0x7FC0_0000). |
| NZF | Nonzero finite number. |
| Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). Can also occur with two nonzero finite number source operands. |
| $Q(x)$ | Return a QNaN with the payload of x . |
| $S(x, y)$ | Return the normalized sum of floating-point value $x$ and negated floating-point value $y$, having unbounded range and precision. Note: If $x=-y$, $v$ is considered to be an exact-zero-difference result (Rezd). |
| $M(x, y)$ | Return the normalized product of floating-point value $x$ and floating-point value y , having unbounded range and precision. |
| p | The intermediate product having unbounded range and precision. |
| $v$ | The intermediate result having unbounded range and precision. |

Table 134.Actions for xvnmsub(alm)sp

## VSX Vector Round to Double-Precision Integer using round to Nearest Away xX2-form

$$
\text { xvrdpi } \quad \mathrm{XT}, \mathrm{XB}
$$



| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XB | $\leftarrow$ BX \\|B |
| ex_flag $\quad \leftarrow$ | $\leftarrow 0 b 0$ |

$$
\text { if }(\text { ex_flag }=0) \text { then VSR }[X T] \leftarrow \text { result }
$$

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let src be the double-precision floating-point operand in doubleword element $i$ of $\operatorname{VSR}[X B]$.
src is rounded to an integer using the rounding mode Round to Nearest Away.

The result is placed into doubleword element $i$ of VSR[XT] in double-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

FX VXSNAN

VSR Data Layout for xvrdpi
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| DP | DP |
| :---: | :---: |

tgt $=$ VSR[XT]

|  | DP |
| :--- | :--- |
| 0 | 64 |

## VSX Vector Round to Double-Precision Integer Exact using Current rounding mode XX2-form

$$
\text { xvrdpic } \quad \mathrm{XT}, \mathrm{XB}
$$



```
XT }\leftarrowTX|
XB
ex_flag \leftarrow0 0b0
do i=0 to 127 by 64
    reset_xflags()
    src{0:63} & VSR[XB]{i:i+63}
    if(RN=0b00) then
        result{i:i+63} }\leftarrow\mathrm{ RoundToDPIntegerNearEven(src)
    if(RN=0b01) then
            result{i:i+63} \leftarrow RoundToDPIntegerTrunc(src)
    if(RN=0b10) then
            result{i:i+63} \leftarrow RoundToDPIntegerCeil(src)
    if(RN=0b11) then
        result{i:i+63} < RoundToDPIntegerFloor(src)
    if(vxsnan_flag) then SetFX(vXSNAN)
    if(xx_flag) then SetFX(XX)
    ex_flag }\leftarrow\mathrm{ ex_flag | (VE & vxsnan_flag)
    ex_flag \leftarrowex_flag|(XE & Xx_flag)
end
if( ex_flag = 0) then VSR[XT]}\leftarrow result
```

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let src be the double-precision floating-point operand in doubleword element $i$ of $\operatorname{VSR}[X B]$.
src is rounded to an integer using the rounding mode specified by RN.

The result is placed into doubleword element i of VSR[XT] in double-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[XT].

## Special Registers Altered

FX XX VXSNAN

VSR Data Layout for xvrdpic
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| DP | DP |
| ---: | :--- |
| tgt = VSR[XT] |  |
| DP | DP |
| 0 | 64 |

VSX Vector Round to Double-Precision Integer using round toward -Infinity XX2-form

| xvrdpim XT, XB |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll}  & 60 \\ 0 \end{array}$ | ${ }_{6} \quad \mathrm{~T}$ | 111 II | 16 B | 21 | 249 | $\left\lvert\, \begin{aligned} & \text { BX TX } \\ & 30 \\ & 31 \\ & \end{aligned}\right.$ |
| XT $\quad \leftarrow T X \\| T$ |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
| ex_flag $\leftarrow 0 \mathrm{ob0}$ |  |  |  |  |  |  |
| do i=0 to 127 by 64 |  |  |  |  |  |  |
| reset_xflags() |  |  |  |  |  |  |
| result $\{\mathrm{i}: 1+63\} \leftarrow$ RoundToDPIntegerFloor (VSR[XB]\{i:i+63\}) |  |  |  |  |  |  |
| if(vxsnan_flag) then SetFX(VXSNAN) |  |  |  |  |  |  |
| ex_flag $\leftarrow$ ex_flag \| (VE \& vxsnan_flag) |  |  |  |  |  |  |
| end |  |  |  |  |  |  |
| if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result |  |  |  |  |  |  |

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element ifrom 0 to 1, do the following. Let $\operatorname{src}$ be the double-precision floating-point operand in doubleword element $i$ of VSR[XB].
src is rounded to an integer using the rounding mode Round toward -Infinity.

The result is placed into doubleword element i of VSR[ XT] in double-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered <br> FX VXSNAN

VSR Data Layout for xvrdpim
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| DP | DP |
| :---: | :---: |

tgt $=$ VSR[XT]

| DP | DP |
| :--- | :--- |
| 0 | 64 |

## VSX Vector Round to Double-Precision Integer using round toward +Infinity XX2-form



| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

$$
\text { do } \mathrm{i}=0 \text { to } 127 \text { by } 64
$$

reset_xflags()

$$
\text { result }\{i: i+63\} \leftarrow \text { RoundToDPIntegerCeil(VSR[XB] }\{i: 1+63\})
$$

if(vxsnan_flag) then SetFX(VXSNAN)
end

$$
\text { ex_flag } \quad \leftarrow \text { ex_flag | (VE \& vxsnan_flag) }
$$

if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1, do the following. Let $\operatorname{src}$ be the double-precision floating-point operand in doubleword element $i$ of VSR[ XB].
$\operatorname{src}$ is rounded to an integer using the rounding mode Round toward +Infinity.

The result is placed into doubleword element i of VSR[ XT] in double-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

Special Registers Altered
FX VXSNAN

## VSR Data Layout for xvrdpip

$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| DP | DP |
| :---: | :---: |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

| DP | DP |
| :--- | :--- |
| 0 | 64 |

## VSX Vector Round to Double-Precision Integer using round toward Zero XX2-form

| xvrdpiz XT,XB |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll}  & 60 \\ 0 & \end{array}$ | ${ }_{6} \quad \mathrm{~T}$ | ${ }_{11}$ III | 16 | 21 | 217 | $\left\lvert\, \begin{aligned} & \text { BX } \\ & 30 \mid \\ & 30 \\ & 31\end{aligned}\right.$ |
| XT $\quad \leftarrow \mathrm{TX} \\|$ T |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
| ex_flag $\leftarrow 0 \mathrm{~b} 0$ |  |  |  |  |  |  |
| do i=0 to 127 by 64 |  |  |  |  |  |  |
| reset_xflags() |  |  |  |  |  |  |
| result $\{i: 1+63\} \leftarrow$ RoundToDPIntegerTrunc(VSR[XB]\{i:i+63\}) |  |  |  |  |  |  |
| if(vxsnan_flag) then SetFX(VXSNAN) |  |  |  |  |  |  |
| ex_flag $\quad \leftarrow$ ex_flag \| (VE \& vxsnan_flag) |  |  |  |  |  |  |
| end |  |  |  |  |  |  |
| if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result |  |  |  |  |  |  |

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element $i$ from 0 to 1 , do the following. Let $\operatorname{src}$ be the double-precision floating-point operand in doubleword element $i$ of VSR[ XB].
src is rounded to an integer using the rounding mode Round toward Zero.

The result is placed into doubleword element i of VSR[XT] in double-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered

FX VXSNAN

VSR Data Layout for xvrdpiz
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| DP | DP |
| :---: | :---: |

$\operatorname{tgt}=\mathrm{VSR}[\mathrm{XT}]$

|  | DP |
| :--- | :--- |
| 0 | 64 |

## Version 3.0

## VSX Vector Reciprocal Estimate Double-Precision XX2-form

$$
\text { xvredp } \quad \text { XT,XB }
$$



| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0$ b0 |

$$
\begin{aligned}
& \text { do } \mathrm{i}=0 \text { to } 127 \text { by } 64 \\
& \text { reset_xflags() }
\end{aligned}
$$

v\{0:inf\} $\leftarrow$ ReciprocalEstimateDP(VSR[XB]\{i:i+63\})
result $\{\mathrm{i}: \mathrm{i}+63\} \leftarrow$ RoundToDP(RN, v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(zx_flag) then SetFX(ZX)
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag | (OE \& ox_flag)
ex_flag $\leftarrow$ ex_flag $\mid$ (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (ZE \& zx_flag)
end
if( ex_flag $=0$ ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let src be the double-precision floating-point operand in doubleword element $i$ of VSR[ XB].

A double-precision floating-point estimate of the reciprocal of src is placed into doubleword element i of VSR[ XT] in double-precision format.

Unless the reciprocal of src would be a zero, an infinity, or a QNaN, the estimate has a relative error in precision no greater than one part in 16384 of the reciprocal of src . That is,

$$
\left|\frac{\text { estimate }-\frac{1}{s r c}}{\frac{1}{s r c}}\right| \leq \frac{1}{16384}
$$

Operation with various special values of the operand is summarized below.

| Source Value | Result | Exception |
| :---: | :---: | :---: |
| -Infinity | -Zero | None |
| -Zero | -Infinity $^{1}$ | ZX |
| +Zero | +lnfinity $^{1}$ | ZX |
| +Infinity | +Zero | None |
| SNaN | QNaN $^{2}$ | VXSNAN |
| QNaN | QNaN | None |

1. No result if $Z E=1$.
2. No result if $\mathrm{VE}=1$.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

The results of executing this instruction is permitted to vary between implementations, and between different executions on the same implementation.

## Special Registers Altered

FX OX UX ZX VXSNAN

VSR Data Layout for xvredp
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| DP | DP |
| :---: | :---: |

tgt $=$ VSR[XT]

| DP | DP |
| :--- | :--- |
| 0 | 64 |

## VSX Vector Reciprocal Estimate Single-Precision XX2-form

| xvresp XT,XB |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll}  & 60 \\ 0 \end{array}$ | ${ }_{6} \quad \mathrm{~T}$ | ${ }_{11} \text { III }$ | $1{ }_{16}$ B | 21 | 154 | $\|$BX TX <br> 30  <br> 31  |
| XT $\quad \leftarrow T X \\| T$ |  |  |  |  |  |  |
| $X B \quad \leftarrow B X$ |  |  |  |  |  |  |
| ex_flag $\leftarrow$ 0bo |  |  |  |  |  |  |
| do i=0 to 127 by 32 |  |  |  |  |  |  |
| reset_xflags() |  |  |  |  |  |  |
| v 0: inf $\} \quad \leftarrow$ ReciprocalEstimateSP(VSR[XB]\{i:i+31\}) |  |  |  |  |  |  |
| result $\{\mathrm{i}: 1+31\} \leftarrow$ RoundToSP(RN, v) |  |  |  |  |  |  |
| if(vxsnan_flag) then SetFX(VXSNAN) |  |  |  |  |  |  |
| if(ox_flag) then SetFX(OX) |  |  |  |  |  |  |
| if(ux_flag) then SetFX(UX) |  |  |  |  |  |  |
| if(zx_flag) then SetFX(ZX) |  |  |  |  |  |  |
| ex_flag $\leftarrow$ ex_flag \| (VE \& vxsnan_flag) |  |  |  |  |  |  |
| ex_flag $\leftarrow$ ex_flag \| (0E \& ox_flag) |  |  |  |  |  |  |
| ex_flag $\leftarrow$ ex_flag \| (UE \& ux_flag) |  |  |  |  |  |  |
| ex_flag $\leftarrow$ ex_flag \| (ZE \& zx_flag) |  |  |  |  |  |  |
| end |  |  |  |  |  |  |
| if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result |  |  |  |  |  |  |

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element $i$ from 0 to 3 , do the following. Let src be the single-precision floating-point operand in word element i of VSR[ XB].

A single-precision floating-point estimate of the reciprocal of src is placed into word element i of VSR[ XT] in single-precision format.

Unless the reciprocal of src would be a zero, an infinity, or a QNaN, the estimate has a relative error in precision no greater than one part in 16384 of the reciprocal of src . That is,

$$
\left|\frac{\text { estimate }-\frac{1}{\mathrm{src}}}{\frac{1}{\mathrm{src}}}\right| \leq \frac{1}{16384}
$$

Operation with various special values of the operand is summarized below.

| Source Value | Result | Exception |
| :---: | :---: | :---: |
| -Infinity | -Zero | None |
| -Zero | -Infinity ${ }^{1}$ | ZX |
| +Zero | +Infinity ${ }^{1}$ | ZX |
| +Infinity | +Zero | None |
| SNaN | QNaN ${ }^{2}$ | VXSNAN |
| QNaN | QNaN | None |

1. No result if $\mathrm{ZE}=1$.
2. No result if $\mathrm{VE}=1$.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

The results of executing this instruction is permitted to vary between implementations, and between different executions on the same implementation.

## Special Registers Altered

FX OX UX ZX VXSNAN

VSR Data Layout for xvresp
src = VSR[ XB]

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

tgt $=$ VSR[ XT]

| SP | SP | SP | SP |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :---: | :---: |
| 0 | 32 | 64 |  |  |  |  |

VSX Vector Round to Single-Precision Integer using round to Nearest Away XX2-form

| xvrspi XT,XB |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll}  & 60 \\ 0 \end{array}$ | ${ }_{6} \quad \mathrm{~T}$ | $\int_{11} \quad \text { III }$ | 16 B | 21 | 137 | $\|$BX TX  <br> 30 31 |
| XT $\quad \leftarrow$ TX \\|| T |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
| ex_flag $\leftarrow 0 \mathrm{bb}$ |  |  |  |  |  |  |
| do i=0 to 127 by 32 |  |  |  |  |  |  |
| reset_xflags() |  |  |  |  |  |  |
| resul | $i: i+31\}$ an_flag | RoundTos | Integer | result $\{1: 1+31\} \leftarrow$ RoundToSPIntegerNearAway (VSR[XB] $] i: 1+31\}$ ) | [ XB ] |  |
| ex_flag $\leftarrow$ ex_flag \| (VE \& vxsnan_flag) |  |  |  |  |  |  |
| end |  |  |  |  |  |  |
| if( ex_flag = 0 ) then VSR $[\mathrm{XT}] \leftarrow$ result |  |  |  |  |  |  |

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3 , do the following. Let src be the single-precision floating-point operand in word element i of $V$ SR[ XB].
src is rounded to an integer using the rounding mode Round to Nearest Away.

The result is placed into word element $i$ of $V S R[X T]$ in single-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered

FX VXSNAN

## VSR Data Layout for xvrspi

$\operatorname{SrC}=$ VSR[ XB]

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

## tgt = VSR[XT]

| SP | SP | SP |  | SP |  |
| :--- | :--- | :--- | :--- | ---: | :---: |
| 0 | 32 | 64 | 96 | 127 |  |

VSX Vector Round to Single-Precision Integer Exact using Current rounding mode XX2-form

| xvrspic $\quad$ XT, XB |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll} 60 \\ 0 \end{array}$ | ${ }_{6} \mathrm{~T}$ | ${ }_{11} \quad \text { III }$ |  | B | 21 | 171 | BX 3 \|TX 3031 |


| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

do i=0 to 127 by 32
reset_xflags()
$\operatorname{src}\{0: 31\} \leftarrow \operatorname{VSR}[X B]\{i: i+31\}$
if( $\mathrm{RN}=0 \mathrm{~b} 00$ ) then
result $\{\mathrm{i}: \mathrm{i}+31\} \leftarrow$ RoundToSPIntegerNearEven(src)
if ( $\mathrm{RN}=0 \mathrm{~b} 01$ ) then
result $\{\mathrm{i}: \mathrm{i}+31\} \leftarrow$ RoundToSPIntegerTrunc(src)
if (RN=0b10) then
result $\{\mathrm{i}: 1+31\} \leftarrow$ RoundToSPIntegerCeil(src)
if ( $\mathrm{RN}=0 \mathrm{b11}$ ) then
result $\{\mathrm{i}: 1+31\} \leftarrow$ RoundToSPIntegerFloor(src)
if(vxsnan_flag) then SetFX(VXSNAN)
if(xx_flag) then $\operatorname{SetFX}(X X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
end
if( ex_flag = 0) then VSR $[\mathrm{XT}] \leftarrow$ result
Let XT be the value $32 \times T X+\mathrm{T}$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 3 , do the following. Let src be the single-precision floating-point operand in word element $i$ of VSR[ XB].
src is rounded to an integer value using the rounding mode specified by RN.

The result is placed into word element $i$ of VSR[ XT] in single-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered

FX XX VXSNAN

## VSR Data Layout for xvrspic

$\operatorname{src}=\operatorname{VSR[XB]}$

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

tgt $=$ VSR[ $X T]$

| SP | SP | SP |  | SP |  |
| :--- | :--- | :--- | :--- | :--- | :---: |
| 0 | 32 | 64 | 96 | 127 |  |

VSX Vector Round to Single-Precision Integer using round toward -Infinity XX2-form


Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element $i$ from 0 to 3 , do the following. Let $\operatorname{src}$ be the single-precision floating-point operand in word element $i$ of VSR[ XB].
src is rounded to an integer using the rounding mode Round toward -Infinity.

The result is placed into word element $i$ of VSR[ XT] in single-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered <br> FX VXSNAN

VSR Data Layout for xvrspim
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

tgt $=\mathrm{VSR}[\mathrm{XT}]$

| SP | SP |  | SP |  |
| :--- | :--- | :--- | :--- | :--- |
| 0 | SP |  |  |  |

VSX Vector Round to Single-Precision Integer using round toward +Infinity XX2-form
xvrspip $\quad \mathrm{XT}, \mathrm{XB}$


| XT | $\leftarrow \mathrm{TX} \\| \mathrm{T}$ |
| :--- | :--- |
| XB | $\leftarrow \mathrm{BX} \\| \mathrm{B}$ |
| ex_flag | $\leftarrow$ 0b0 |

do i=0 to 127 by 32
reset_xflags()
result $\{i: i+31\}=$ RoundToSPIntegerCeil(VSR[XB]\{i:i+31\})
if(vxsnan_flag) then SetFX(VXSNAN)
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
end
if( ex_flag = 0) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element $i$ from 0 to 3 , do the following. Let $\operatorname{src}$ be the single-precision floating-point operand in word element $i$ of VSR[ XB].
src is rounded to an integer using the rounding mode Round toward +Infinity.

The result is placed into word element $i$ of VSR[ XT] in single-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

```
Special Registers Altered
FX VXSNAN
```

VSR Data Layout for xvrspip
$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

$\operatorname{tgt}=\mathrm{VSR}[\mathrm{XT}]$

| SP | SP |  | SP |  |
| :--- | :--- | :--- | :--- | :--- |

VSX Vector Round to Single-Precision Integer using round toward Zero XX2-form


Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element $i$ from 0 to 3 , do the following. Let $\operatorname{src}$ be the single-precision floating-point operand in word element i of VSR[ XB].
src is rounded to an integer using the rounding mode Round toward Zero.

The result is placed into word element $i$ of VSR[ XT ] in single-precision format.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered

FX VXSNAN

## VSR Data Layout for xurspiz

$\operatorname{src}=$ VSR[XB]

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

$$
\text { tgt }=\operatorname{VSR}[X T]
$$

| SP | SP |  | SP |  |
| :--- | :--- | :--- | :--- | :--- |

VSX Vector Reciprocal Square Root Estimate Double-Precision XX2-form

| xvrsqrte |  | XT,XB |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll}  & 60 \\ 0 & \end{array}$ | $6$ | ${ }_{11} \quad \text { III }$ | 16 | 21 | 202 | BX  <br> 30 31 |
| XT $\quad \leftarrow \mathrm{TX} \\|$ T |  |  |  |  |  |  |
| $X B \quad \leftarrow B X$ |  |  |  |  |  |  |
| ex_flag $\leftarrow 0 \mathrm{~b} 0$ |  |  |  |  |  |  |
| do $\mathrm{i} \leftarrow 0$ to 127 by 64 |  |  |  |  |  |  |
| reset_xflags() |  |  |  |  |  |  |
| $v\{0: i n f\} \quad \leftarrow$ RecipSquareRootEstimateDP(VSR[XB]\{i:i+63\}) |  |  |  |  |  |  |
| result $\{\mathrm{i}: 1+63\} \leqslant$ RoundToDP(RN, v) |  |  |  |  |  |  |
| if(vxsnan_flag) then SetFX(VXSNAN) |  |  |  |  |  |  |
| if(vxsqrt_flag) then SetFX(VXSQRT) |  |  |  |  |  |  |
| if(zx_flag) then SetFX(ZX) |  |  |  |  |  |  |
| ex_flag $\leftarrow$ ex_flag \| (VE \& vxsnan_flag) |  |  |  |  |  |  |
| ex_flag |  | ex_flag | (VE \& vxsqrt_flag) |  |  |  |
| ex_flag |  | ex_flag | (ZE \& zx_flag) |  |  |  |
| end |  |  |  |  |  |  |
| if( ex_flag = 0 ) then VSR[XT] $\leftarrow$ result |  |  |  |  |  |  |

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1, do the following. Let src be the double-precision floating-point operand in doubleword element $i$ of VSR[ XB].

A double-precision floating-point estimate of the reciprocal square root of SrC is placed into doubleword element i of VSR[XT] in double-precision format.

Unless the reciprocal of the square root of src would be a zero, an infinity, or a QNaN, the estimate has a relative error in precision no greater than one part in 16384 of the reciprocal of the square root of src . That is,

$$
\left|\frac{\text { estimate }-\frac{1}{\sqrt{s r c}}}{\frac{1}{\sqrt{\mathrm{src}}}}\right| \leq \frac{1}{16384}
$$

Operation with various special values of the operand is summarized below.

| Source Value | Result | Exception |
| :---: | :---: | :---: |
| -Infinity | QNaN $^{1}$ | VXSQRT |
| +Infinity | +Zero | None |
| -Finite | QNaN $^{1}$ | VXSQRT |
| -Zero | -Infinity $^{2}$ | ZX |
| +Zero | +lnfinity $^{2}$ | ZX |
| SNaN | QNaN $^{1}$ | VXSNAN |
| QNaN | QNaN | None |

1. No result if $\mathrm{VE}=1$
2. No result if $\mathrm{ZE}=1$.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

The results of executing this instruction is permitted to vary between implementations, and between different executions on the same implementation.

## Special Registers Altered

> FX ZX VXSNAN VXSQRT

VSR Data Layout for xvrsqrtedp
src = VSR[XB]

| DP | DP |
| :---: | :---: |

tgt $=$ VSR[ XT]

|  | DP |
| :--- | :--- |
| 0 | 64 |

## Version 3.0

VSX Vector Reciprocal Square Root Estimate Single-Precision XX2-form
xvrsqrtesp $\quad \mathrm{XT}, \mathrm{XB}$

| -60 | 6 | T | 11 | III | 16 | B | 21 | 138 | TX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

```
XT 
XB }
ex_flag \leftarrow0b0
do i=0 to 127 by 32
    reset_xflags()
    v{0:inf}}\leftarrow<\mathrm{ RecipSquareRootEstimateSP(VSR[XB]{i:i+31})
    result{i:i+31} \leftarrow RoundToDP(RN,v)
    if(vxsnan_flag) then SetFX(VXSNAN)
    if(vxsqrt_flag) then SetFX(vXSQRT)
    if(zx_flag) then SetFX(ZX)
    ex_flag \leftarrowex_flag| (VE & vxsnan_flag)
    ex_flag \leftarrowex_flag| (VE & vxsqrt_flag)
    ex_flag }\leftarrow\mathrm{ ex_flag | (ZE & Zx_flag)
end
if( ex_flag = 0 ) then VSR[XT] \leftarrow result
```

Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element ifrom 0 to 3 , do the following. Let src be the single-precision floating-point operand in word element $i$ of VSR[ XB].

A single-precision floating-point estimate of the reciprocal square root of src is placed into word element i of VSR[ XT] in single-precision format.

Unless the reciprocal of the square root of src would be a zero, an infinity, or a QNaN, the estimate has a relative error in precision no greater than one part in 16384 of the reciprocal of the square root of src . That is,

$$
\left|\frac{\text { estimate }-\frac{1}{\sqrt{\mathrm{src}}}}{\frac{1}{\sqrt{\mathrm{src}}}}\right| \leq \frac{1}{16384}
$$

Operation with various special values of the operand is summarized below.

| Source Value | Result | Exception |
| :---: | :---: | :---: |
| -Infinity | QNaN $^{1}$ | VXSQRT |
| +Infinity | +Zero | None |
| -Finite | QNaN $^{1}$ | VXSQRT |
| -Zero | -Infinity |  |
| +Zero | +lnfinity $^{2}$ | ZX |
| SNaN | QNaN $^{1}$ | VXSNAN |
| QNaN | QNaN | None |

1. No result if $\mathrm{VE}=1$.
2. No result if $\mathrm{ZE}=1$.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

The results of executing this instruction is permitted to vary between implementations, and between different executions on the same implementation.

Special Registers Altered
FX ZX VXSNAN VXSQRT

## VSR Data Layout for xvrsqrtesp

$\mathrm{src}=\mathrm{VSR}[\mathrm{XB}]$

| $S P$ | $S P$ | $S P$ | $S P$ |
| :---: | :---: | :---: | :---: |

$$
\operatorname{tgt}=\operatorname{VSR}[\mathrm{XT}]
$$

| SP | SP |  | SP |
| :--- | :--- | :--- | :--- |

## VSX Vector Square Root Double-Precision XX2-form

$$
\text { xvsqrtdp } \quad \mathrm{XT}, \mathrm{XB}
$$



| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

do $i \leftarrow 0$ to 127 by 64
reset_xflags()
$\mathrm{v}\{0:$ inf $\} \quad \leftarrow$ SquareRootDP(VSR[XB]\{i:i+63\})
result $\{\mathrm{i}: i+63\} \leftarrow$ RoundToDP (RN, v)
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxsqrt_flag) then SetFX(VXSQRT)
if( $x x$ _flag) then $\operatorname{SetFX}(X X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vxsqrt_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag
end
if( ex_flag ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element $i$ from 0 to 1 , do the following. Let $\operatorname{src}$ be the double-precision floating-point operand in doubleword element $i$ of VSR[ XB].

The unbounded-precision square root of $\operatorname{src}$ is produced.

See Table 135.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

| src |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\begin{aligned} & \hline \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \text { vxsqrt_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline v \leftarrow d Q N a N \\ & \text { vxsqri_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{V} \leftarrow+$ Zero | $\mathrm{V} \leftarrow+$ Zero | $v \leftarrow$ SQRT(src) | $V \leftarrow+$ Infinity | $\mathrm{V} \leftarrow \mathrm{SrC}$ | $\begin{aligned} & \hline \mathrm{v} \leftarrow Q(\mathrm{src}) \\ & \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{aligned}$ |


| Explanation: |  |
| :--- | :--- |
| src | The double-precision floating-point value in doubleword element $i$ of VSR[XB] (where $i \in\{0,1\})$. |
| dQNaN | Default quiet $\operatorname{NaN}\left(0 \times 7 F F 8 \_0000 \_0000 \_0000\right)$. |
| NZF | Nonzero finite number. |
| SQRT $(x)$ | The unbounded-precision square root of the floating-point value $x$. |
| $Q(x)$ | Return a QNaN with the payload of $x$. |
| $v$ | The intermediate result having unbounded signficand precision and unbounded exponent range. |

Table 135.Actions for xvsqrtdp

## Version 3.0

## VSX Vector Square Root Single-Precision XX2-form

## xvsqrtsp XT,XB



| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XB | $\leftarrow$ BX \\| B |
| ex_flag | $\leftarrow 0 b 0$ |

$$
\text { do } i=0 \text { to } 127 \text { by } 32
$$

reset_xflags()

$$
\text { v\{0:inf }\} \quad \leftarrow \text { SquareRootSP(VSR[XB]\{i:i+31\}) }
$$

$$
\text { result }\{\dot{1}: i+31\} \leftarrow \text { RoundToSP(RN, v) }
$$

if(vxsnan_flag) then SetFX(VXSNAN)
if(vxsqrt_flag) then SetFX(VXSQRT)

$$
\text { if(xx_flag) then } \operatorname{SetFX}(X X)
$$

$$
\text { ex_flag } \quad \leftarrow \text { ex_flag | (VE \& vxsnan_flag) }
$$

$$
\text { ex_flag } \leftarrow \text { ex_flag } \mid \text { (VE \& vxsqrt_flag) }
$$

end

$$
\text { ex_flag } \leftarrow \text { ex_flag | (XE \& xx_flag }
$$

if( ex_flag ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element $i$ from 0 to 3 , do the following. Let $\operatorname{src}$ be the single-precision floating-point operand in word element $i$ of VSR[ XB].

The unbounded-precision square root of src is produced.

See Table 136.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

| src |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\mathrm{V} \leftarrow \mathrm{dQNaN}$ vxsart_flag $\leftarrow 1$ | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \mathrm{vxsqq} \text { _lag } \leftarrow 1 \end{aligned}$ | $v \leftarrow+$ Zero | $v \leftarrow+$ Zero | $\mathrm{v} \leqslant$ SQRT(src) | $v \leftarrow+$ lnfinity | V \& SrC | $\begin{aligned} & \hline \begin{array}{l} \mathrm{v} \leftarrow Q(\text { src }) \\ \text { vxsnan_flag } \leftarrow 1 \end{array} \end{aligned}$ |


| Explanation: |  |
| :--- | :--- |
| src | The single-precision floating-point value in word element i of $\operatorname{VSR}[\mathrm{XB}]$ (where $\mathrm{i} \in\{0,1,2,3\}$ ). |
| dQNaN | Default quiet NaN (0x7FC0_0000). |
| NZF | Nonzero finite number. |
| SQRT $(\mathrm{x})$ | The unbounded-precision square root of the floating-point value x. |
| Q x . | Return a QNaN with the payload of x. |
| V | The intermediate result having unbounded signficand precision and unbounded exponent range. |

Table 136.Actions for xvsqrtsp

## VSX Vector Subtract Double-Precision XX3-form

$$
\text { xvsubdp } \quad \text { XT,XA,XB }
$$

| 60 |  | T |  | A |  | B |  | 104 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 |  | 6 |  | 11 |  | 16 |  | 21 |


| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

do $\mathrm{i}=0$ to 127 by 64
reset_xflags()
$\operatorname{src} 1 \quad \leftarrow \operatorname{VSR}[X A]\{i: i+63\}$
src2 $\leftarrow$ VSR[XB] $\{i: i+63\}$
v\{0:inf\} $\quad \leftarrow \operatorname{AddDP}(\operatorname{src} 1$, NegateDP(src2))
result $\{i: i+63\} \leftarrow \operatorname{RoundTODP}(R N, v)$
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxisi_flag) then SetEX(VXISI)
if(ox_flag) then SetFX(0X)
if(ux_flag) then $\operatorname{SetFX}(U X)$
if(xx_flag) then SetFX(XX)
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vxisi_flag)
ex_flag $\leftarrow$ ex_flag | ( 0 \& \& ox_flag)
ex_flag $\leftarrow$ ex_flag | (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
end
if( ex_flag ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element i from 0 to 1 , do the following. Let srcl be the double-precision floating-point operand in doubleword element $i$ of VSR[ XA].

Let srcz be the double-precision floating-point operand in doubleword element $i$ of VSR[ XB].
$\operatorname{src}$ is negated and added ${ }^{[1]}$ to srcl, producing a sum having unbounded range and precision.

The sum is normalized ${ }^{[2]}$.
See Table 137.
The intermediate result is rounded to double-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

[^10]|  |  | src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\overline{\mathrm{O}}$ | -Infinity | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \mathrm{vxisis} \text { flag } \leftarrow 1 \end{aligned}$ | v ¢-Infinity | v - Infinity | $v \leftarrow-$ Infinity | $v \leftarrow-$ Infinity | $\mathrm{v} \leftarrow-$ Infinity | v \& Src2 | $\begin{aligned} & \hline \begin{array}{l} \mathrm{v} \leftarrow Q(\text { src } 2) \\ \mathrm{vxsnn} \text { _flag } \leftarrow 1 \end{array} \end{aligned}$ |
|  | -NZF | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leftarrow \mathrm{S}$ (src1, src2) | $\mathrm{v} \leqslant \mathrm{src} 1$ | $\mathrm{v} \leqslant \mathrm{src} 1$ | $\mathrm{v} \leqslant \mathrm{S}$ (src1, src2) | $v \leftarrow-$ Infinity | v < Src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srce2) } \\ & \mathrm{vxsnan} \mathrm{flag} \leftarrow 1 \end{aligned}$ |
|  | -Zero | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leftarrow-\mathrm{Sr} \mathrm{C}^{2}$ | $\mathrm{v} \leftarrow$-Zero | $v \leftarrow$ Rezd | $\mathrm{v} \leftarrow-\mathrm{SrC2}$ | $\mathrm{v} \leftarrow-$ Infinity | v ¢ Src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src } 2) \\ & \mathrm{vxsnn} \text { _flag } \leftarrow 1 \end{aligned}$ |
|  | +Zero | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leftarrow-\mathrm{Sr} \mathrm{C}^{2}$ | $v \leftarrow$ Rezd | $v \leftarrow+$ Zero | $v \leftarrow-\mathrm{Sr} 2$ | $\mathrm{v} \leftarrow-$ Infinity | v \& Src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { ssco2) } \\ & \mathrm{vxsnan} \mathrm{flag} \leftarrow 1 \end{aligned}$ |
|  | +NZF | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leqslant \mathrm{S}$ (src1, src2) | $\mathrm{v} \leqslant \mathrm{src} 1$ | $\mathrm{v} \leqslant$ SrCl | $v \leftarrow S($ src1, src2) | $v \leftarrow-$ Infinity | v ¢ Src2 | $\begin{aligned} & \hline \begin{array}{l} v \leftarrow Q(s r c 2) \\ \text { vxsnan_flag } \leftarrow 1 \end{array} \\ & \hline \end{aligned}$ |
|  | +Infinity | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{V} \leftarrow+$ lnininity | $v \leftarrow+$ Infinity | $\mathrm{V} \leftarrow+$ lnininity | $v \leftarrow+$ lnfinity | $\mathrm{v} \leftarrow \mathrm{dQNaN}$ <br> vxisi_flag $\leftarrow 1$ | v < Src2 | $\begin{aligned} & v \leftarrow Q(\text { src2 } 2) \\ & v x s n a n \_f l a g \leftarrow 1 \end{aligned}$ |
|  | QNaN | $\mathrm{v} \leqslant \mathrm{src} 1$ | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant$ Src 1 | $v \leftarrow$ SrC1 | $v \leftarrow$ SrC 1 | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant \mathrm{src} 1$ | $\begin{aligned} & \mathrm{v} \leftarrow \text { Src1 } \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |
|  | SNaN | $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{srcc}) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{V} \leftarrow \mathrm{Q}$ (src1) vxsnan_flag $\leftarrow 1$ | $\begin{array}{\|l\|} \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src } 1) \\ \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{array}$ | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { Ilag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline v \leftarrow Q(\text { srcl }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline \begin{array}{l} \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{srcc}) \\ \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{array} \end{aligned}$ | $\begin{aligned} & \hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |

[^11]Table 137.Actions for xvsubdp

## VSX Vector Subtract Single-Precision XX3-form

$$
\text { xvsubsp } \quad \text { XT,XA,XB }
$$



| XT | $\leftarrow T X \\| T$ |
| :--- | :--- |
| XA | $\leftarrow A X \\| A$ |
| XB | $\leftarrow B X \\| B$ |
| ex_flag | $\leftarrow 0 b 0$ |

do $i=0$ to 127 by 32
reset_xflags()
$\operatorname{src1} \leftarrow \operatorname{VSR}[X A]\{i: i+31\}$
$\operatorname{src2} \quad \leftarrow \operatorname{VSR}[X B]\{i: i+31\}$
v\{0:inf $\} \quad \leftarrow$ AddSP(src1,NegateSP(src2))
result $\{\mathrm{i}: i+31\} \leftarrow$ RoundToSP $(R N, v)$
if(vxsnan_flag) then SetFX(VXSNAN)
if(vxisi_flag) then SetFX(VXISI)
if(ox_flag) then SetFX(OX)
if(ux_flag) then SetFX(UX)
if(xx_flag) then $\operatorname{SetFX}(X X)$
ex_flag $\leftarrow$ ex_flag | (VE \& vxsnan_flag)
ex_flag $\leftarrow$ ex_flag | (VE \& vxisi_flag)
ex_flag $\quad \leftarrow$ ex_flag | (OE \& ox_flag)
ex_flag $\leftarrow$ ex_flag | (UE \& ux_flag)
ex_flag $\leftarrow$ ex_flag | (XE \& xx_flag)
end
if( ex_flag ) then VSR[XT] $\leftarrow$ result
Let $X T$ be the value $32 \times T X+T$.
Let $X A$ be the value $32 \times A X+A$.
Let $X B$ be the value $32 \times B X+B$.
For each vector element $i$ from 0 to 3 , do the following. Let srcl be the single-precision floating-point operand in word element $i$ of VSR[ XA].

Let srcl be the single-precision floating-point operand in word element $i$ of VSR[ XB].
$\operatorname{src}$ is negated and added ${ }^{[1]}$ to srcl, producing a sum having unbounded range and precision.

The sum is normalized ${ }^{[2]}$.
See Table 138.
The intermediate result is rounded to single-precision using the rounding mode specified by RN.

See Table 58, "Scalar Floating-Point Intermediate Result Handling," on page 516.

The result is placed into word element $i$ of VSR[ XT] in single-precision format.

See Table 106, "Vector Floating-Point Final Result," on page 663.

If a trap-enabled exception occurs in any element of the vector, no results are written to VSR[ XT] .

## Special Registers Altered

FX OX UX XX VXSNAN VXISI

VSR Data Layout for xvsubsp
src1 = VSR[XA]

| SP | SP | SP | SP |
| :---: | :---: | :---: | :---: |
| src2 = VSR[XB] |  |  |  |
| SP | SP | SP | SP |

tgt $=$ VSR $[\mathrm{XT}]$

| SP | SP | SP | SP |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :---: | :---: |
| 0 | 32 | 64 |  |  |  |  |

[^12]|  |  | src2 |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | -Infinity | -NZF | -Zero | +Zero | +NZF | +Infinity | QNaN | SNaN |
| $\overline{\mathrm{O}}$ | -Infinity | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{dQNaN} \\ & \mathrm{vxisis} \text { flag } \leftarrow 1 \end{aligned}$ | v ¢-Infinity | v - Infinity | $v \leftarrow-$ Infinity | $v \leftarrow-$ Infinity | $\mathrm{v} \leftarrow-$ Infinity | v \& Src2 | $\begin{aligned} & \hline \begin{array}{l} \mathrm{v} \leftarrow Q(\text { src } 2) \\ \mathrm{vxsnn} \text { _flag } \leftarrow 1 \end{array} \end{aligned}$ |
|  | -NZF | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leftarrow \mathrm{S}$ (src1, src2) | $\mathrm{v} \leqslant \mathrm{src} 1$ | $\mathrm{v} \leqslant \mathrm{src} 1$ | $\mathrm{v} \leqslant \mathrm{S}$ (src1, src2) | $v \leftarrow-$ Infinity | v < Src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { srce2) } \\ & \mathrm{vxsnan} \mathrm{flag} \leftarrow 1 \end{aligned}$ |
|  | -Zero | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leftarrow-\mathrm{Sr} \mathrm{C}^{2}$ | $\mathrm{v} \leftarrow$-Zero | $v \leftarrow$ Rezd | $\mathrm{v} \leftarrow-\mathrm{SrC2}$ | $\mathrm{v} \leftarrow-$ Infinity | v ¢ Src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { src } 2) \\ & \mathrm{vxsnn} \text { _flag } \leftarrow 1 \end{aligned}$ |
|  | +Zero | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leftarrow-\mathrm{Sr} \mathrm{C}^{2}$ | $v \leftarrow$ Rezd | $v \leftarrow+$ Zero | $v \leftarrow-\mathrm{Sr} 2$ | $\mathrm{v} \leftarrow-$ Infinity | v \& Src2 | $\begin{aligned} & \mathrm{v} \leftarrow Q(\text { ssco2) } \\ & \mathrm{vxsnan} \mathrm{flag} \leftarrow 1 \end{aligned}$ |
|  | +NZF | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{v} \leqslant \mathrm{S}$ (src1, src2) | $\mathrm{v} \leqslant \mathrm{src} 1$ | $\mathrm{v} \leqslant$ SrCl | $v \leftarrow S($ src1, src2) | $v \leftarrow-$ Infinity | v ¢ Src2 | $\begin{aligned} & \hline \begin{array}{l} v \leftarrow Q(s r c 2) \\ \text { vxsnan_flag } \leftarrow 1 \end{array} \\ & \hline \end{aligned}$ |
|  | +Infinity | $\mathrm{v} \leftarrow+$ Infinity | $\mathrm{V} \leftarrow+$ lnininity | $v \leftarrow+$ Infinity | $\mathrm{V} \leftarrow+$ lnininity | $v \leftarrow+$ lnfinity | $\mathrm{v} \leftarrow \mathrm{dQNaN}$ <br> vxisi_flag $\leftarrow 1$ | v < Src2 | $\begin{aligned} & v \leftarrow Q(\text { src2 } 2) \\ & v x s n a n \_f l a g \leftarrow 1 \end{aligned}$ |
|  | QNaN | $\mathrm{v} \leqslant \mathrm{src} 1$ | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant$ Src 1 | $v \leftarrow$ SrC1 | $v \leftarrow$ SrC 1 | $\mathrm{v} \leqslant$ SrC1 | $\mathrm{v} \leqslant \mathrm{src} 1$ | $\begin{aligned} & \mathrm{v} \leftarrow \text { Src1 } \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |
|  | SNaN | $\begin{aligned} & \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{srcc}) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\mathrm{V} \leftarrow \mathrm{Q}$ (src1) vxsnan_flag $\leftarrow 1$ | $\begin{array}{\|l\|} \hline \mathrm{v} \leftarrow \mathrm{Q}(\text { src } 1) \\ \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{array}$ | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { Ilag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline v \leftarrow Q(\text { srcl }) \\ & \text { vxsnan_flag } \leftarrow 1 \end{aligned}$ | $\begin{aligned} & \hline \begin{array}{l} \mathrm{v} \leftarrow \mathrm{Q}(\mathrm{srcc}) \\ \mathrm{vxsnan} \text { _flag } \leftarrow 1 \end{array} \end{aligned}$ | $\begin{aligned} & \hline \mathrm{V} \leftarrow \mathrm{Q}(\mathrm{src} 1) \\ & \mathrm{vxsnan} \text { flag } \leftarrow 1 \end{aligned}$ |

[^13]\leftarrowVVR[XA]{i:i+31
src2 }\leftarrowV\operatorname{VR[XB]{i:i+31}
e_a }\leftarrow\mathrm{ src1{1:8} - 127
e_b }\leftarrow\operatorname{src2{1:8}-127
fe_flag \leftarrowfe_flag | IsNaN(src1) | IsInf(src1) |
IsNaN(src2) | IsInf(src2) | IsZero(src2) |
(e_b <= -126 ) |
(e_b>= 125)|
( !IsZero(src1) \& ( (e_a - e_b) >= 127 ) ) |
( !IsZero(src1) \& ( (e_a - e_b) <= -125 ) ) |
( !IsZero(src1) \& ( e_a <= -103 ) )
fg_flag \leftarrowfg_flag | IsInf(src1) | IsInf(src2) |
IsZero(src2) | IsDen(src2)
end
fl_flag }\leftarrow\mathrm{ xvredp_error() <= 2-14
CR[BF]}\leftarrow0b1 || fg_flag || fe_flag || 0b

```

Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
\(f e_{-} f l a g\) is initialized to 0 .
\(f_{g_{-}} f l a g\) is initialized to 0 .
For each vector element \(i\) from 0 to 3 , do the following. Let srcl be the single-precision floating-point operand in word element \(i\) of \(V S R[X A]\).

Let \(\mathrm{srcl}_{2}\) be the single-precision floating-point operand in word element i of \(V\) SR[ XB].

Let \(e_{-}\)a be the unbiased exponent of srcl.
Let \(e_{-} b\) be the unbiased exponent of \(\operatorname{src}\).
fe_flag is set to 1 for any of the following conditions.
- srcl is a NaN or an infinity.
- srcl is a zero, a NaN, or an infinity.
\(-e_{-} b\) is less than or equal to -126 .
\(-e_{-}^{-} b\) is greater than or equal to 125.
- sricl is not a zero and the difference, \(e_{-} a \cdot e_{-} b\), is greater than or equal to 127 .
- srcl is not a zero and the difference, \(e_{-} a \cdot e_{-} b\), is less than or equal to -125 .
- sricl is not a zero and e_a is less than or equal to-103.
fg_flag is set to 1 for any of the following conditions.
- srcl is an infinity.
- src2 is a zero, an infinity, or a denormalized value.
\(C R\) field \(B F\) is set to the value Ob1 ||fg_flag ||fe_flag || Obo.

\section*{Special Registers Altered}

CR[BF]
VSR Data Layout for xvtdivsp
srcl = VSR[ XA]
\begin{tabular}{|l|l|l|l|}
\hline .word[0] & .word[1] & .word[2] & .word[3] \\
\hline
\end{tabular}
\(\operatorname{srC2}=\) VSR[ \(X B]\)


\section*{VSX Vector Test for software Square Root Double-Precision XX2-form}
\[
\begin{aligned}
& \text { xvtsqrtdp BF,XB } \\
& \text { end } \\
& \text { fl_flag } \leftarrow \text { xvrsqrtedp_error () }<=2^{-14} \\
& C R[B F] \leftarrow 0 b 1 \text { || fg_flag || fe_flag || 0b0 }
\end{aligned}
\]

Let \(X B\) be the value \(32 \times B X+B\).
fe fl ag is initialized to 0 .
\(\mathrm{fg}_{\mathrm{f}} \mathrm{fl} \mathrm{ag}\) is initialized to 0 .
For each vector element \(i\) from 0 to 1 , do the following. Let src be the double-precision floating-point operand in doubleword element \(i\) of VSR[XB].

Let \(\mathrm{e} \_\mathrm{b}\) be the unbiased exponent of src .
fe_flag is set to 1 for any of the following conditions.
- \(\operatorname{src}\) is a zero, a NaN, an infinity, or a negative value.
\(-e_{-} b\) is less than or equal to 970 .
\(f_{g_{-}} f \mid a g\) is set to 1 for the following condition.
- src is a zero, an infinity, or a denormalized value.
\(C R\) field \(B F\) is set to the value Ob1 ||fg_fag ||fe_fag ||Obo.

\section*{Special Registers Altered}

CR[BF]

\section*{VSR Data Layout for xvtsqrtdp}
\(\operatorname{src}=\) VSR[ XB]
\begin{tabular}{|l|l|}
\hline. dword[0] &.\(d\) word[1] \\
\hline \(0 \quad\) & 64 \\
\hline
\end{tabular}

\section*{VSX Vector Test for software Square Root Single-Precision XX2-form}
xvtsqrtsp BF,XB
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 60 & BF & II & & III & & B & \\
\hline
\end{tabular}
\[
X B \quad \leftarrow B X \| B
\]
\[
\text { fe_flag } \leftarrow 0 b 0
\]
\[
\text { fg_flag } \leftarrow 0 b 0
\]
do \(\mathrm{i}=0\) to 127 by 32
\[
\text { src } \quad \leftarrow \operatorname{VSR}[X B]\{i: i+31\}
\]
\[
\text { e_b } \leftarrow \operatorname{src} 2\{1: 8\}-127
\]
\[
\text { fe_flag } \leftarrow \text { fe_flag }|\operatorname{ISNaN(src)}| \operatorname{IsInf(src)~|}
\]
IsZero(src) | IsNeg(src) |(e_a <=-103)
\[
\text { fg_flag } \leftarrow \text { fg_flag }\left.\right|^{\text {IsInf(src) }} \mid \text { Iszero(src) } \mid
\]
IsDen(src)
end
fl_flag \(=\) xvrsqrtesp_error() \(<=2^{-14}\)
\(C R[B F]=0 b 1\) || fg_flag || fe_flag || 0 bo
Let \(X B\) be the value \(32 \times B X+B\).
fe fl ag is initialized to 0 .
\(\mathrm{f}_{\mathrm{g}} \mathrm{fl} \mathrm{ag}\) is initialized to 0 .
For each vector element \(i\) from 0 to 3 , do the following. Let \(\operatorname{src}\) be the single-precision floating-point operand in word element \(i\) of VSR[ XB].

Let \(e_{-} b\) be the unbiased exponent of \(\operatorname{src}\).
fe_flag is set to 1 for any of the following conditions.
- src is a zero, a NaN, an infinity, or a negative value.
\(-e_{-} b\) is less than or equal to 103.
\(\mathrm{fg}_{\mathrm{f}} \mathrm{f} \operatorname{lag}\) is set to 1 for the following condition.
- \(\operatorname{src}\) is a zero, an infinity, or a denormalized value.
\(C R\) field \(B F\) is set to the value Ob1 ||fg_flag ||fe_flag || 0 bo.

\section*{Special Registers Altered}

CR[BF]

VSR Data Layout for xvtsqrtsp
src = VSR[XB]
\begin{tabular}{|l|l|l|l|}
\hline. word[0] &. word [ 1] & .word[2] & .word[3] \\
\hline 0 & \multicolumn{2}{|c|}{\({ }_{32}\)} \\
\hline
\end{tabular}

\section*{VSX Vector Test Data Class Double-Precision XX2-form}
xvtstdcdp XT,XB,DCMX
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \[
60
\] & \[
T
\] & \[
\int_{11} d x
\] & \[
{ }_{16} \quad B
\] & \[
15
\] & \({ }_{\text {dc }}^{\text {dc }}\) 25 \(2^{\text {26 }} 5\) &  \\
\hline
\end{tabular}
if MSR. VSX=0 then VSX_Unavailable ()
\(D C M x \leftarrow d c\|d m\| d x\)
do \(i=0\) to 1
\begin{tabular}{|c|c|}
\hline sic & \(\leftarrow\) VSR[ \(32 \times 8 \times+B]\). dword[i] \\
\hline sign & \(\leftarrow\) src.bitlol \\
\hline exponent & \(\leftarrow\) src. bit [1:11] \\
\hline fraction & \(\leftarrow\) src.bit [12:63] \\
\hline
\end{tabular}
class.Infinity \(\leftarrow(\) exponent \(=\) Ox7FF) \& (fraction \(=0)\)
class. NaN \(\leftarrow(\) exponent \(=0 \times 7\) FF) \& (fraction \(!=0)\)
class.Zero \(\leftarrow(\) exponent \(=\) oxooo \() \&(\) fraction \(=0)\)
class. Denormal \(\leftarrow(\) exponent \(=0 x 000) \&(\) fraction \(!=0)\)
match \(\leftarrow(\) DCMX.bitiol \& class. NaN \()\)
(DCHX, bit[1] \& class.Infinity \& !sign)
(DCMX, bit[2] \& class.Infinity \& sign)
(DCMX, bit[3] \& class.Zero \& !sign)
(DCMX. bit[4] \& class.Zero \& sign)
( DCMX. bit[5] \& class. Denormal \& \(\operatorname{sign}\) )
(DCMX, bit[6] \& class. Denormal \& sign)
if match \(=1\) then
VSR[32xTXXT] , dword[i] \(\leftarrow\) OxFFFF_FFFF_FFFF_FFFF
else
VSR[32xTX+T].dword[i] 1 Ox0000_0000_0000_0000
end

Let \(X B\) be the sum \(32 \times B X+B\).
Let XT be the sum \(32 \times T X+\mathrm{T}\).
Let DCMX be the value \(d c\) concatenated with \(d x\) concatenated with dcm .

For each integer value i from 0 to 1 , do the following. Let \(\operatorname{src}\) be the double-precision floating-point value in doubleword element \(i\) of \(\operatorname{VSR}[\mathrm{XB}]\).

If src matches one of the 7 possible data classes specified by DCMX (Data Class Mask), the contents of doubleword element i of VSR[XT] are set to \(\mathrm{OXFFFF}_{2}\) FFF_FFFFFFFF . Otherwise, the contents of doubleword element i of VSR[XT] are set to \(0 \times 0000 \_0000 \_0000 \_0000\).
\begin{tabular}{cl} 
DCMX bit & Data Class \\
0 & NaN \\
1 & +Infinity \\
2 & - Infinity \\
3 & +Zero \\
4 & - Zero \\
5 & +Denormal \\
6 & - Denormal
\end{tabular}

Special Registers Altered: None

VSR Data Layout for xvtstdcdp
\begin{tabular}{|c|c|c|c|}
\hline sic & VSR[ XB]. dword[0] & VSR[ XB ]. dword[1] & \\
\hline tgt & VSR[ XT] . dword[0] & VSR[ XT].dword[1] & \\
\hline & 0 & 64 & 127 \\
\hline
\end{tabular}

VSX Vector Test Data Class Single-Precision XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|l|}{xvtstdcsp \(\mathrm{XT}, \mathrm{XB}, \mathrm{DCMX}\)} \\
\hline \[
0
\] & \[
{ }_{6} \quad \mathrm{~T}
\] & \(\int_{11} \mathrm{dx}\) & 16 B & \({ }_{21} 13\) & \(\left\lvert\,\)\begin{tabular}{l|l|l|} 
dc \\
25
\end{tabular}\({ }_{26} 5\right.\) & dmm
\(29 \times 1\)
30 \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablell
DCMX \(\leftarrow d c\|d m\| d x\)
do \(i=0\) to 3
SrC \(\leftarrow \operatorname{VSR}[32 \times B X+B]\). word \([i]\)
sign \(\quad \leftarrow\) sic.bit \([0]\)
exponent \(\leftarrow\) src.bit[1:8]
fraction \(\leftarrow\) src.bit[9:31]
class.Infinity \(\leftarrow(\) exponent \(=\) OxFF \() \&(\) fraction \(=0)\)
class. NaN \(\leftarrow(\) exponent \(=0 \times F F) \&(f r a c t i o n!=0)\)
class.Zero \(\leftarrow(\) exponent \(=\) OxOO \() \&(f r a c t i o n=0)\)
class. Denormal \(\leftarrow(\) exponent \(=0 \times 00) \&(f r a c t i o n!=0)\)
match \(\leftarrow(\) DCMX.bitl0] \& class. NaN \()\) (DCMX, bit[1] \& class.Infinity \& !sign) (DCMX, bit[2] \& class.Infinity \& sign) (DCMX, bit[3] \& class.Zero \& !sign) (DCMX, bit[4] \& class.Zero \& sign) (DCMX, bit[5] \& class. Denormal \& !sign) | (DCMX, bitl6] \& class. Denormal \& sign)

Let \(X B\) be the sum \(32 \times B X+B\).
Let \(X T\) be the sum \(32 \times T X+T\).
Let DCMX be the value \(d c\) concatenated with \(d m\) concatenated with \(d x\).

For each integer value i from 0 to 3 , do the following. Let \(s r_{c}\) be the single-precision floating-point value in word element \(i\) of VSR[XB].

If src matches one of the 7 possible data classes specified by DCMX (Data Class Mask), the contents of word element i of VSR[XT] are set to OXFFFF_FFFF. Otherwise, the contents of word element \(i\) of VSR[ XT] are set to \(0 \times 0000 \_0000\).
\begin{tabular}{cl} 
DCMX bit & Data Class \\
0 & NaN \\
1 & +Infinity \\
2 & - Infinity \\
3 & +Zero \\
4 & - Zero \\
5 & +Denormal \\
6 & - Denormal
\end{tabular}

\section*{Special Registers Altered:}

None
    if match \(=1\) then
        VSR[32xTX+T].dword[i] \(\quad\) XxFFFF_FFFF
    else
        VSR[32xTXTT].dword [ \(]\) ] \(0 \times 0000\) _0000
end
| VSR Data Layout for xvtstdcsp
\begin{tabular}{c|c|c|c|c|}
\hline \multirow{3}{*}{\(\operatorname{src}\)} & VSR[XB].word[0] & VSR[XB].word[1] & VSR[XB].word[2] & VSR[XB].word[3] \\
\hline
\end{tabular}
tgt
\begin{tabular}{|l|l|l|l|}
\hline VSR[XT].word[0] & VSR[XT].word[ 1] & VSR[XT].word[2] & VSR[XT].word[3] \\
\hline 0 & 32 & 64 & 96
\end{tabular}

\section*{VSX Vector Extract Exponent Double-Precision XX2-form}
xvxexpdp XT,XB
\begin{tabular}{|l|ll|l|l|l|l|l|}
\hline 60 & & T & & 0 & & B & \\
\hline 0 & & 6 & & 11 & & 16 & \\
21 & & & 375 & \\
30031 \\
\hline
\end{tabular}
if MSR. VSX=O then VSX_Unavailable()
do \(\mathrm{i}=0\) to 1
sic \(\leftarrow\) VSR \([32 \times B X+B]\). dword \([i]\)
VSR[32xTX+T]. dword[i] \(\leftarrow\) EXTZ64(sic. bit[1:11])
end
Let \(X T\) be the sum \(32 \times T X+T\).
Let \(X B\) be the sum \(32 \times B X+B\).
For each integer value i from 0 to 1, do the following.
Let \(\operatorname{src}\) be the double-precision floating-point value in doubleword element \(i\) of VSR[ XB].

The value of the exponent field in src is placed into doubleword element \(i\) of VSR[ XT] in unsigned integer format.

Special Registers Altered:
None

VSX Vector Extract Exponent Single-Precision XX2-form
\begin{tabular}{|c|c|c|c|c|c|}
\hline xvxexp & & T,XB & & & \\
\hline \[
\begin{array}{ll} 
& 60 \\
0 & \\
\hline
\end{array}
\] & \[
6
\] & \[
{ }_{11} 8
\] & \[
16
\] & \[
24
\] &  \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailable()
do \(\mathrm{i}=0\) to 3
\(\operatorname{src} \leftarrow V S R[32 \times B X+B]\), word \([i]\)

end
Let \(X T\) be the sum \(32 \times T X+T\).
Let \(X B\) be the sum \(32 \times B X+B\).
For each integer value i from 0 to 3 , do the following. Let src be the single-precision floating-point value in word element i of \(\operatorname{VSR[XB]}\).

The value of the exponent field in src is placed into word element \(i\) of VSR[ XT] in unsigned integer format.

Special Registers Altered: None

\section*{VSR Data Layout for xvxexpdp}


\section*{VSR Data Layout for xvxexpsp}
\(\square\)
\begin{tabular}{c|l|l|l|l|}
\hline tgt & VSR[XT].word[0] & VSR[XT].word[1] & VSR[XT].word[2] & VSR[XT].word[3] \\
\hline 0 & 32 & 64 & 96
\end{tabular}

\section*{VSX Vector Extract Significand Double-Precision XX2-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{xvxsigdp XT,XB} \\
\hline \[
60
\] & \[
6
\] & \[
11
\] & \[
16
\] & \[
21 \quad 475
\] & \(\left\lvert\, \begin{aligned} & \text { BX } \\ & 30 \mid \\ & 31 \\ & 31\end{aligned}\right.\) \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailable(l)
do \(i=0\) to 1
Sr \(c \leftarrow \operatorname{VSR}[32 \times B X+B]\).dword \([i]\)
exponent \(\leftarrow\) EXTZ(src.bit[1:11])
fraction \(\leftarrow\) EXT264(src. bit[12:63])
if (exponent ! = 0) \& (exponent \(!=2047\) ) then
fraction \(\leftarrow\) fraction \(\mid 0 \times 0010_{-} 0000 \_0000 \_0000\)
VSR[ \(32 \times T X+T]\). dword \([i] \leftarrow\) fraction
end

Let \(X T\) be the sum \(32 \times T X+T\).
Let \(X B\) be the sum \(32 \times B X+B\).
For each integer value i from 0 to 1 , do the following. Let \(\operatorname{src}\) be the double-precision floating-point value in doubleword element \(i\) of VSR[ XB].

The significand of src is placed into doubleword element \(i\) of VSR[ XT] in unsigned integer format. If src is a normal value, the implicit leading bit is set to 1.

Special Registers Altered:
None

VSX Vector Extract Significand Single-Precision XX2-form

\section*{xvxsigsp \\ XT,XB}
\begin{tabular}{|l|ll|ll|l|ll|l|}
\hline 60 & \multicolumn{2}{|c|}{T} & \multicolumn{2}{c|}{9} & & B & & 475 \\
0 & & 6 & & 11 & & 16 & & 21 \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailable(l)
do \(i=0\) to 3
src \(\leftarrow \operatorname{VSR}[32 \times B X+B]\), word \([i]\)
exponent \(\leftarrow\) EXTZ(src.bit[1:8])
fraction \(\leftarrow\) EXTZ32(src. bit[9:31])
if (exponent \(!=0\) ) \& (exponent \(!=255\) ) then
fraction \(\leftarrow\) fraction \(\mid\) Ox0080_0000
\(\operatorname{VSR}[32 \times T X+T]\) word \([i] \leftarrow\) fraction
end
Let \(X T\) be the sum \(32 \times T X+T\).
Let \(X B\) be the sum \(32 \times B X+B\).
For each integer value i from 0 to 3 , do the following. Let src be the single-precision floating-point value in word element \(i\) of VSR[ XB].

The significand of \(s r c\) is placed into word element i of VSR[XT] in unsigned integer format. If \(s r c\) is a normal value, the implicit leading bit is set to 1 .

Special Registers Altered:
None

VSR Data Layout for xvxsigdp
\begin{tabular}{|c|c|c|c|}
\hline sic & VSR[ XB].dword[0] & VSR[ XB]. dword[1] & \\
\hline \multirow[t]{2}{*}{tgt} & VSR[ XT].dword[0] & VSR[ XT]. dword[1] & \\
\hline & 0 & 64 & 127 \\
\hline
\end{tabular}

\section*{VSR Data Layout for xvxsigsp}
\begin{tabular}{|c|c|c|c|}
\hline VSR[ XB]. word[0] & VSR[ XB]. word[ 1] & VSR[ XB]. word[ 2] & VSR[ XB]. word[3] \\
\hline
\end{tabular}
tgt \begin{tabular}{|l|l|l|l|}
\hline VSR[XT].word[0] & VSR[ XT]. word[1] & VSR[XT]. word[2] & VSR[XT].word[3] \\
\hline 0 & 32 & 64 & 96
\end{tabular}

\section*{VSX Vector Byte-Reverse Doubleword XX2-form}

if MSR.VSX=0 then VSX_Unavailable(l)
do \(\mathrm{i}=0\) to 1
do \(j=0\) to 7
\(\operatorname{VSR}[32 \times T X+T]\). dword[i], byte[j] \(\leftarrow \operatorname{VSR}[32 \times B X+B]\). dword[i], byte[7-j] end
end
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each integer value \(i\) from 0 to 1 , do the following. The contents of byte 7 of doubleword element \(i\) of VSR[XB] are placed into byte 0 of doubleword element i of \(\operatorname{VSR}[\mathrm{XT}]\).

The contents of byte 6 of doubleword element \(i\) of VSR[XB] are placed into byte 1 of doubleword element i of \(\operatorname{VSR[XT].}\)

The contents of byte 5 of doubleword element \(i\) of VSR[XB] are placed into byte 2 of doubleword element i of VSR[ XT].

The contents of byte 4 of doubleword element \(i\) of VSR[XB] are placed into byte 3 of doubleword element I of VSR[ XT].

The contents of byte 3 of doubleword element \(i\) of VSR[XB] are placed into byte 4 of doubleword element i of VSR[ XT].

The contents of byte 2 of doubleword element \(i\) of VSR[XB] are placed into byte 5 of doubleword element i of VSR[ XT].

The contents of byte 1 of doubleword element \(i\) of VSR[XB] are placed into byte 6 of doubleword element i of VSR[ XT].

The contents of byte 0 of doubleword element \(i\) of VSR[XB] are placed into byte 7 of doubleword element i of VSR[ XT].

\section*{Special Registers Altered: \\ None}

\section*{VSX Vector Byte-Reverse Halfword XX2-form}
xxbrh XT,XB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|}
\hline & 60 & & T & & 7 & & B & & 475 & BXITX \\
\hline 0 & & 6 & & 11 & & 16 & & 21 & & 3031 \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailable()
do \(i=0\) to 7
VSR[32xTX + T], hword[i], byte \([0] \leftarrow \operatorname{VSR}[32 \times B X+B]\), hword[i], byte[1]
\(\operatorname{VSR}[32 \times T X+T]\), hword[i], byte[1] \(\operatorname{VSR}[32 \times B X+B]\), hword[i], byte[0]
end
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each integer value i from 0 to 7 , do the following. The contents of byte 1 of halfword element \(i\) of VSR[ XB] are placed into byte 0 of halfword element i of VSR[XT].

The contents of byte 0 of halfword element \(i\) of VSR[ XB] are placed into byte 1 of halfword element i of VSR[XT].

\section*{Special Registers Altered: None}

VSX Vector Byte-Reverse Quadword XX2-form
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline \multicolumn{7}{|l|}{xxbrq \(\quad\) XT, XB} \\
\hline \[
60
\] & \({ }_{6}{ }^{\text {T }}\) & \[
{ }_{11} 31
\] & 16 & 21 & 475 & \(\left\lvert\, \begin{aligned} & \text { Bx } \\ & 30 \times 1 \times 1 \\ & 31\end{aligned}\right.\) \\
\hline
\end{tabular}
if MSR. VSXX=0 then VSX_Unavailable()
do \(i=0\) to 15
VSR[ \(32 x T X+T]\), byte \([i] \leftarrow \operatorname{VSR}[32 \times B X+B]\), byte \([15 \cdot i]\)
end
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each integer value i from 0 to 15 , do the following. The contents of byte sub-element \(15 \cdot \mathrm{i}\) of VSR[ XB] are placed into byte sub-element \(i\) of VSR[ XT].

Special Registers Altered:
None

VSX Vector Byte-Reverse Word XX2-form

if MSR.VSX=O then VSX_Unavailable el)
do \(i=0\) to 3
do \(j=0\) to 3 \(\operatorname{VSR}[32 x T X+T]\), vord \([i]\), byte \([j] \leftarrow \operatorname{VSR}[32 x B X+B]\), word \([i]\), byte \([3 \cdot j]\) end
end
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
For each integer value i from 0 to 3 , do the following. The contents of byte 3 of word element \(i\) of VSR[ XB] are placed into byte 0 of word element i of VSR[ XT].

The contents of byte 2 of word element i of VSR[ XB] are placed into byte 1 of word element i of VSR[ XT].

The contents of byte 1 of word element i of VSR[ XB] are placed into byte 2 of word element i of VSR[ XT].

The contents of byte 0 of word element i of VSR[ XB] are placed into byte 3 of word element i of VSR[ XT].

Special Registers Altered:
None

VSX Vector Extract Unsigned Word XX2-form xxextractuw \(\quad\) XT,XB,UIM

if MSR.VSX=0 then VSX_Unavailablel)

VSR[32xTX+T].dword[0] \(\leftarrow\) Chop \((\) EXTZ(sic \(), 64)\)
VSR[32xTX+T], dword[1] © Ox0000_0000_0000_0000
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of byte elements UIM:UIM+3 of VSR[XB] are placed into word element 1 of VSR[XT]. The contents of the remaining word elements of VSR[XT] are set to 0 .

If the value of \(U I M\) is greater than 12, the results are undefined.

\section*{Special Registers Altered:}

None

\section*{VSX Vector Insert Word XX2-form}
xxinsertw \(\quad\) XT,XB,UIM
\begin{tabular}{|l|l|l|l|l|l|l|l|l|}
\hline 60 & & T & \multicolumn{1}{|l|}{} & UIM & & B & & 181 \\
11 & 12
\end{tabular}
if MSR. VSX=0 then VSX_Unavailablel)
VSR[32xTX+T], byte[ Ul M: UI W+3] \(\leftarrow\) VSR[ \(32 \times B X+B]\), bit [ \(32: 63]\)
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of word element 1 of VSR[ XB] are placed into byte elements \(\mathrm{UI} \operatorname{M}: \mathrm{Ul} \operatorname{M+3}\) of VSR[XT]. The contents of the remaining byte elements of VSR[XT] are not modified.

If the value of \(U \mid M\) is greater than 12 , the results are undefined.

\section*{Special Registers Altered:}

None

VSX Logical AND XX3-form
xxland
XT, XA, XB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & 11 & A & 16 & B & 21 & 130 &  \\
\hline
\end{tabular}
\begin{tabular}{ll}
\(X T\) & \(\leftarrow T X \| T\) \\
\(X A\) & \(\leftarrow A X \| A\) \\
\(X B\) & \(\leftarrow B X \| B\)
\end{tabular}
\(\operatorname{VSR}[X T] \leftarrow \operatorname{VSR}[X A] \& \operatorname{VSR}[X B]\)
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of VSR[XA] are ANDed with the contents of VSR[ XB] and the result is placed into VSR[ XT].


VSX Logical AND with Complement XX3-form
xxlandc \(\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & 11 & A & 16 & B & 21 & 138 & AXXBXTX 293031 \\
\hline
\end{tabular}
\begin{tabular}{ll} 
XT & \(\leftarrow T X \| T\) \\
XA & \(\leftarrow A X \| A\) \\
\(X B\) & \(\leftarrow B X \| B\) \\
VSR[XT] & \(\leftarrow \operatorname{VSR}[X A] \& \sim \operatorname{VSR}[X B]\)
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of VSR[XA] are ANDed with the complement of the contents of VSR[ XB] and the result is placed into VSR[ XT].

\section*{Special Registers Altered \\ None}

VSR Data Layout for xxland
srcl = VSR[XA]
\(\square\)
SIC2 \(=\) VSR[ XB]

tgt \(=\) VSR[ XT]
\begin{tabular}{|l|}
\hline \\
0
\end{tabular}

\section*{Version 3.0}

\section*{VSX Logical Equivalence XX3-form}

\[
\operatorname{VSR}[32 \times T X+T] \leftarrow \operatorname{VSR}[32 \times A X+A] \equiv \operatorname{VSR}[32 \times B X+B]
\]

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of VSR[XA] are exclusive-ORed with the contents of VSR[XB] and the complemented result is placed into VSR[ XT].

\section*{Special Registers Altered:}

None

\section*{VSR Data Layout for xxleqv}
\(\operatorname{SrC}=\) VSR[ XA]

src = VSR[XB]

tgt \(=\) VSR[ XT]


\section*{VSX Logical NAND XX3-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{10}{|l|}{xxInand \(\quad X T, X A, X B\)} \\
\hline \[
\begin{array}{ll}
60 \\
\hline
\end{array}
\] & 6 & T & 11 & A & 16 & B & 21 & 178 & AXX \({ }_{2} \times\) PTX
293031 \\
\hline
\end{tabular}
\(\operatorname{VSR}[32 \times T X+T] \leftarrow \neg(\operatorname{VSR}[32 \times A X+A] \& \operatorname{VSR}[32 \times B X+B])\)
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of VSR[XA] are ANDed with the contents of \(V S R[X B]\) and the complemented result is placed into VSR[ XT].

\section*{Special Registers Altered:}

\section*{None}

\section*{VSR Data Layout for xxInand}

SIC = VSR[ XA]

tgt \(=\) VSR[ XT]
\(\square\)

VSX Logical OR with Complement XX3-form
xxlorc
XT, XA, XB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & 11 & A & 16 & B & 21 & 170 &  \\
\hline
\end{tabular}
\(\operatorname{VSR}[32 \times T X+T] \leftarrow \operatorname{VSR}[32 \times A X+A] \mid\) VSR \([32 \times B X+B]\)
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

The contents of VSR[XA] are ORed with the complement of the contents of VSR[XB] and the result is placed into VSR[ XT].

Special Registers Altered:
None
VSR Data Layout for xxlorc
srcl = VSR[ XA]

tgt \(=\) VSR[ \(X T]\)


VSX Logical NOR XX3-form
xxInor \(\quad X T, X A, X B\)
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & 11 & A & 16 & B & 21 & 162 & \[
|a x B x| T x
\] \\
\hline
\end{tabular}
\(\operatorname{VSR}[32 \times T X+T] \leftarrow \sim(\operatorname{VSR}[32 \times A X+A] \mid \operatorname{VSR}[32 \times B X+B])\)
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

The contents of VSR[ XA] are ORed with the contents of VSR[XB] and the complemented result is placed into VSR[XT].

\section*{Special Registers Altered}

None

VSR Data Layout for xxInor
srcl = VSR[XA]

\(\operatorname{srC2}=\operatorname{VSR}[X B]\)

tgt = VSR[ XT]


\section*{Version 3.0}

\section*{VSX Logical OR XX3-form}

\(\operatorname{VSR}[32 \times T X+T] \leftarrow \operatorname{VSR}[32 \times A X+A] \mid \operatorname{VSR}[32 \times B X+B]\)
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of VSR[XA] are ORed with the contents of VSR[ \(X B]\) and the result is placed into VSR[ \(X T]\).


\section*{VSX Logical XOR XX3-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{10}{|l|}{xxlxor \(\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}\)} \\
\hline 060 & 6 & & 11 & A & 16 & B & 21 & 154 & \(\left|\begin{array}{l}\text { AXBX } \\ 2930 \times X \\ 31\end{array}\right|\) \\
\hline
\end{tabular}

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of VSR[XA] are exclusive-ORed with the contents of VSR[XB] and the result is placed into VSR[ XT].


\section*{VSX Merge High Word XX3-form}
xxmrghw \(\quad X T, X A, X B\)

if MSR. VSX=O then VSX_Unavailable()
VSR[ \(32 \times T X+T]\). word \([0] \leftarrow \operatorname{VSR}[32 \times A X+A]\). wor d[ 0\(]\)
VSR [32xTX+T]. word [1] \(\leftarrow\) VSR [ \(32 \times B X+B]\). wor d[ 0\(]\)
VSR [32xTXXT]. word [ 2\(] \leftarrow \operatorname{VSR}[32 x A X+A]\). word \([1]\)
\(\operatorname{VSR}[32 x T X+T]\). word \([3] \leftarrow \operatorname{VSR}[32 x B X+B]\). wor d [1]
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of word element 0 of VSR[ XA] are placed into word element 0 of VS R[ XT].

The contents of word element 0 of VSR[ XB] are placed into word element 1 of VSR[ XT] .

The contents of word element 1 of VSR[ XA] are placed into word element 2 of VSR[XT].

The contents of word element 1 of VSR[ XB] are placed into word element 3 of VS R[ XT] .

\section*{Special Registers Altered \\ None}

VSR Data Layout for xxmrghw
srcl = VSR[ XA]
\begin{tabular}{|c|c|c|c|}
\hline .word[0] & . word[1] & unused & unused \\
\hline src2 \(=\operatorname{VSR}[X B]\) \\
\begin{tabular}{|c|c|c|c|}
\hline .word[0] & .word[1] & unused & unused \\
\hline
\end{tabular}
\end{tabular}.
tgt \(=\) VSR[XT]


VSX Merge Low Word XX3-form
xxmrglw \(\quad X T, X A, X B\)
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline 60 & 6 & T & 11 & A & 16 & B & 21 & 50 & \[
|a x B x| T x
\] \\
\hline
\end{tabular}
if MSR.VSX=O then VSX_Unavailablell
\(\operatorname{VSR}[32 x T X+T]\), word \([0] \leftarrow \operatorname{VSR}[32 x A X+A]\), word [ 2\(]\)
VSR [ \(32 \times T X+T]\), word \([1] \leftarrow \operatorname{VSR}[32 \times B X+B]\). word [ 2\(]\)
VSR[ \(32 x T X+T]\), word \([2] \leftarrow \operatorname{VSR}[32 x A X+A]\). word \([3]\)
VSR[ \(32 x T X+T]\), word \([3] \leqslant \operatorname{VSR}[32 x B X+B]\). word [ 3\(]\)
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

The contents of word element 2 of VSR[XA] are placed into word element 0 of VSR[ XT].

The contents of word element 2 of VSR[ XB] are placed into word element 1 of VSR[ XT] .

The contents of word element 3 of VSR[ XA] are placed into word element 2 of VSR[ XT].

The contents of word element 3 of VSR[ XB] are placed into word element 3 of VSR[ XT].

\section*{Special Registers Altered \\ None}

VSR Data Layout for xxmrglw
srcl = VSR[XA]
\begin{tabular}{|c|c|c|c|}
\hline unused & unused & .word[2] & .word[3] \\
\hline srC2 \(=\) VSR[ XB] \\
\begin{tabular}{|c|c|c|c|}
\hline unused & unused & .word [2] & .word [3] \\
\hline
\end{tabular}
\end{tabular}\(.\)\begin{tabular}{l}
\end{tabular}
tgt \(=\) VSR[ XT]
\begin{tabular}{|l|l|l|l|}
\hline. word[0] & .word[1] & .word[2] & .word[3] \\
\hline 0 & 32 & 64 & 96
\end{tabular}

\section*{VSX Vector Permute XX3-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{xxperm \(\quad\) XT, XA. XB} \\
\hline \[
0
\] & \[
{ }_{6} \quad \mathrm{~T}
\] & \({ }_{11}\) & & B & 21 & 26 & \[
\begin{aligned}
& \text { AXBXXTX } \\
& 293031
\end{aligned}
\] \\
\hline
\end{tabular}

> if MSR. VSX=O then VSX_Unavailable()
> sic.byte[0:15] \(\leftarrow\) VSR[ \(32 \times A X+A]\)
> sce. byte[ \(16: 31] \leftarrow \operatorname{VSR}[32 \times T X+T]\)
> pcr.byte[ \(0: 15] \leftarrow\) VSR[ \(32 \times B X+B]\)
> do \(i=0\) to 15
> idx \(\leftarrow\) pcr. byte[i], bit[3:7]
> VSR[32xTX+T], byte[i] \(\operatorname{src}\). byte[idx]
> end

Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(X T\) be the value \(32 \times T X+T\).
Let bytes \(0: 15\) of \(\operatorname{sic}\) be the contents of VSR[ XA].
Let bytes 16:31 of src be the contents of VSR [ XT].
Let the permute control vector pcv be the contents of VSR[ XB].

For each integer value i from 0 to 15 , do the following.
Let idx be the unsigned integer in bits 3:7 of byte element \(i\) of \(p c v\).

The contents of byte element \(i d x\) of \(s r c\) is placed into byte element i of \(\mathrm{VSR}[\mathrm{XT}]\).

\section*{Special Registers Altered:}

None

\section*{VSX Vector Permute Right-indexed XX3-form}
xxpermr \begin{tabular}{l} 
XT,XA.XB \\
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 60 & & T & & A & B & \\
\hline 0 & & & & & 11 & 16
\end{tabular} \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailablel)
```

src.byte[0:15] \leftarrowVSR[32xAX+A]
src. byte[16:31]\leftarrowVSR[32xTX+T]
pcv.byte[0:15]}\leftarrow\operatorname{VSR[32\timesBX+B]
do i = 0 to 15
idx \& pcr.byte[i], bit[3:7]
VSR[32xTX +T], byte[i] \& src. byte[31-idx]
end

```

Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(X T\) be the value \(32 \times T X+T\).
Let bytes \(0: 15\) of \(\operatorname{src}\) be the contents of VSR[ XA].
Let bytes 16:31 of src be the contents of VSR[ XT].
Let the permute control vector pcv be the contents of VSR[ XB].

For each integer value i from 0 to 15 , do the following. Let \(i d x\) be the unsigned integer in bits 3:7 of byte element \(i\) of \(p c v\).

The contents of byte element \(31 \cdot \mathrm{idx}\) of src is placed into byte element i of VSR[ XT].

\section*{Special Registers Altered: \\ None}

\section*{VSX Permute Doubleword Immediate XX3-form}
xxpermdi \(\quad \mathrm{XT}, \mathrm{XA}, \mathrm{XB}, \mathrm{DM}\)

if MSR. VSX \(=0\) then VSX_Unavallable()
VSR[32xTX+T].dwor d[0] \& VSR[32xAXXA]. dwor dl DM, bit t [0]]
\(\operatorname{VsR}[32 x T X+T]\). dword \([1] \leftarrow \operatorname{VSR}[32 \times B X+B]\). dwor d [ DM. bit [1]]
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
If \(\mathrm{DM} . \operatorname{bit}[0]=0\), the contents of doubleword element 0 of VSR[ XA ] are placed into doubleword element 0 of VSR[XT]. Otherwise the contents of doubleword element 1 of VSR[XA] are placed into doubleword element 0 of VSR[ XT].

If DM.bit[1]=0, the contents of doubleword element 0 of VSR[ XB] are placed into doubleword element 1 of VSR[XT]. Otherwise the contents of doubleword element 1 of VSR[XB] are placed into doubleword element 1 of VSR[ XT].

Special Registers Altered
None
\begin{tabular}{lll}
\hline Extended Mnemonic & \multicolumn{2}{l}{ Equivalent To } \\
\hline xxspltd & \(T, A, O\) & xxpermdi \\
\hline xxspltd \(A, A, O b 00\) \\
\hline xxmrghd & \(T, A, 1\) & xxpermdi \\
\hline T, \(, A, A, O b 11\) \\
\hline xxmrgld & \(T, A, B\) & xxpermdi \\
\hline x, \(A, B, O b 00\) \\
\hline
\end{tabular}

VSR Data Layout for xxpermdi
srcl = VSR[ XA]
\begin{tabular}{|l|l|}
\hline.\(d w o r d[0]\) &.\(d w o r d[1]\) \\
\hline
\end{tabular}
\(\operatorname{srC2}=\) VSR[ XB]
\begin{tabular}{|l|l|}
\hline.\(d w o r d[0]\) &.\(d\) word \([1]\) \\
\hline
\end{tabular}
tgt \(=\) VSR[ XT]
\begin{tabular}{|l|l|}
\hline &.\(d\) word [0] \\
\hline 0 & \(.{ }^{2} \quad . d\) word [ 1] \\
\hline
\end{tabular}

VSX Select XX4-form
xxsel \(\quad X T, X A, X B, X C\)
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|}
\hline \({ }^{1} 60\) & 6 & T & 11 & A & 16 & B & & C & 3
26 &  \\
\hline
\end{tabular}
```

if MSR.VSX=O then VSX_Unavailable|)
do i=0 to 127
if (VSR[32xCX+C), bit[i]=0) then
VSR[32xTX+T], bit[i] \& VSR[32xAX+A], bit[i]
else
VSR[32xTX+T], bit[i] <VSR[32xBX+B], bit[i]
end

```

Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).
Let \(X C\) be the value \(32 \times C X+C\).
For each bit of VSR[XC] that contains the value 0 , the corresponding bit of VSR[XA] is placed into the corresponding bit of VSR[XT]. Otherwise, the corresponding bit of VSR[XB] is placed into the corresponding bit of VS R[ XT].


\section*{Version 3.0}

VSX Shift Left Double by Word Immediate XX3-form
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{10}{|l|}{xxsldwi XT,XA,XB,SHW} \\
\hline 060 & 6 & T & & A & 16 & B &  & &  \\
\hline
\end{tabular}
if MSR.VSX=0 then VSX_Unavailable()
```

source.qword[0]}\leftarrow\operatorname{VSR[32xAX +A]
source.qword[ [1] \leftarrowVSR[32\timesBX+B]
VSR[32xTX+T] \leftarrow source.word[SHW:SHW+3]

```
Let \(X T\) be the value \(32 \times T X+T\).
Let \(X A\) be the value \(32 \times A X+A\).
Let \(X B\) be the value \(32 \times B X+B\).

Let the source vector be the concatenation of the contents of VSR[XA] followed by the contents of VSR[XB]. Words SHW: SHW+3 of the source vector are placed into VSR[ XT].

\section*{Special Registers Altered None}

\section*{VSR Data Layout for xxsidwi}
srcl = VSR[XA]
\begin{tabular}{|l|l|l|l|}
\hline .word[0] & .word[1] & .word[2] & .word[3] \\
\hline
\end{tabular}
src2 \(=\) VSR[ XB]
\begin{tabular}{|l|l|l|l|}
\hline .word[0] & .word[1] & .word[2] & .word[3] \\
\hline
\end{tabular}
\[
\operatorname{tgt}=\operatorname{VSR}[X T]
\]
\begin{tabular}{|l|l|l|l|}
\hline . word[0] &, word[1] & , word[2] & , word[3] \\
\hline 0 & 64 & 96 & 127
\end{tabular}

\section*{VSX Vector Splat Immediate Byte}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{9}{|l|}{xxspltib \(\quad\) XT,IMM8} \\
\hline \[
060
\] & 6 & T & \(1{ }_{11} 0\) & 13 & IMM8 & 21 & 360 & Tx| \\
\hline
\end{tabular}
if \(T X=0\) \& MSR. VSX=O then VSX_Unavailable el)
if \(T X=1\) \& MSR.VEC=0 then Vector_Unavailable()
do \(\mathrm{i}=0\) to 15
VSR[32xTX+T], byte \([i] \leftarrow\) Ul M8
end
Let \(X T\) be the sum \(32 \times T X+T\).
The value I MM8 is copied into each byte element of VSR[ XT].

\section*{Special Registers Altered:}

None

\section*{VSX Splat Word XX2-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|}
\hline \multicolumn{11}{|l|}{xxspltw \(\quad\) XT, XB, UIM} \\
\hline 0 & 60 & 6 & T & \({ }_{11}\) III & 14 & 16 & B & 21 & 164 &  \\
\hline
\end{tabular}
if MSR. VSX=0 then VSX_Unavailablel)
\(\operatorname{VSR}[32 \times T X+T]\), word \([0] \leftarrow \operatorname{VSR}[32 \times B X+B]\), word \([\mathrm{Ul}\) M]
\(\operatorname{VSR}[32 \times T X+T]\), word \([1] \leftarrow \operatorname{VSR}[32 \times B X+B]\), word \([\) Ul M]
VSR[ \(32 x T X+T]\), word \([2] \leftarrow \operatorname{VSR}[32 \times B X+B]\). word \([\) UI M\(]\) \(\operatorname{VSR}[32 x T X+T]\), word \([3] \leftarrow \operatorname{VSR}[32 \times B X+B]\). wor d [ Ul W]

Let XT be the value \(32 \times T X+T\).
Let \(X B\) be the value \(32 \times B X+B\).
The contents of word element UIM of VSR[XB] are replicated in each word element of VSR[ XT] .

\section*{Special Registers Altered None}

\section*{VSR Data Layout for xxspltw}
src = VSR[XB]
\begin{tabular}{|l|l|l|l|}
\hline .word[0] & .word[1] & .word[2] & .word[3] \\
\hline
\end{tabular}
tgt \(=\operatorname{VSR}[X T]\)
\begin{tabular}{|l|l|l|l|}
\hline. word[0] &., word[ 1] &. word[ 2] & . word[ 3] \\
\hline 0 & 32 & 64 & 96
\end{tabular}

\section*{Appendix A. Suggested Floating-Point Models}

\section*{A. 1 Floating-Point Round to Single-Precision Model}

The following describes algorithmically the operation of the Floating Round to Single-Precision instruction.
```

If (FRB)}\mp@subsup{)}{1:11}{<<897 and (FRB) 1:63 > 0 then
Do
If FPSCR
If FPSCR UE = 1 then goto Enabled Exponent Underflow
End
If (FRB)}\mp@subsup{)}{1:11}{}>1150\mathrm{ and (FRB) 1:11 < 2047 then
Do
If FPSCR
If FPSCR OE = 1 then goto Enabled Exponent Overflow
End
If (FRB)}\mp@subsup{)}{1:11}{}>896\mathrm{ and (FRB)}\mp@subsup{)}{1:11}{<<1151 then goto Normal Operand
If (FRB)}\mp@subsup{)}{1:63}{}=0\mathrm{ then goto Zero Operand
If (FRB)}\mp@subsup{)}{1:11}{}=2047\mathrm{ then
Do
If (FRB) 12:63 = 0 then goto Infinity Operand
If (FRB) 12 = 1 then goto QNaN Operand
If (FRB) 12 =0 and (FRB) 13:63 >0 then goto SNaN Operand
End

```

\section*{Disabled Exponent Underflow:}
sign \(\leftarrow(\mathrm{FRB})_{0}\)
If \((F R B)_{1: 11}=0\) then
Do
\(\exp \leftarrow-1022\)
\(\operatorname{frac}_{0: 52} \leftarrow 0 \mathrm{ODO}\) II \((\mathrm{FRB})_{12: 63}\)
End
If \((F R B)_{1: 11}>0\) then
Do
\(\exp \leftarrow(F R B)_{1: 11}-1023\)
\(\operatorname{frac}_{0: 52} \leftarrow 0 \mathrm{Ob} 1\) II \((\mathrm{FRB})_{12: 63}\)
End
Denormalize operand:
\(\mathrm{G}\|\mathrm{R}\| \mathrm{X} \leftarrow 0 \mathrm{~b} 000\)
Do while exp <-126
\(\exp \leftarrow \exp +1\)
\(\operatorname{frac}_{0: 52}\|G\| R\|X \leftarrow 0 b 0\| f r a c_{0: 52}\|G\|(R \mid X)\)
End
FPSCR \(_{U X} \leftarrow\left(\mathrm{frac}_{24: 52}\|G\| R \| X\right)>0\)
Round Single(sign,exp,frac \({ }_{0: 52}, \mathrm{G}, \mathrm{R}, \mathrm{X}\) )
FPSCR \(_{X X} \leftarrow\) FPSCR \(_{X X} \mid\) FPSCR \(_{\text {FI }}\)
If \(\mathrm{frac}_{0: 52}=0\) then
Do
\(\mathrm{FRT}_{0} \leftarrow\) sign
\(\mathrm{FRT}_{1: 63} \leftarrow 0\)

\section*{Version 3.0}
```

        If sign \(=0\) then FPSCR \(_{\text {FPRF }} \leftarrow\) "+ zero"
        If sign \(=1\) then FPSCR \(_{\text {FPRF }} \leftarrow\) "- zero"
        End
    If frac $_{0: 52}>0$ then
Do
If frac $_{0}=1$ then
Do
If sign $=0$ then FPSCR $_{\text {FPRF }} \leftarrow$ " + normal number"
If sign $=1$ then FPSCR $_{\text {FPRF }} \leftarrow$ "- normal number"
End
If $\mathrm{frac}_{0}=0$ then
Do
If sign $=0$ then FPSCR $_{\text {FPRF }} \leftarrow$ " + denormalized number"
If sign $=1$ then FPSCR $_{\text {FPRF }} \leftarrow$ "- denormalized number"
End
Normalize operand:
Do while frac ${ }_{0}=0$
$\exp \leftarrow \exp -1$
$\mathrm{frac}_{0: 52} \leftarrow \mathrm{frac}_{1: 52} \mathrm{II} 0 \mathrm{Ob} 0$
End
$\mathrm{FRT}_{0} \leftarrow$ sign
FRT $_{1: 11} \leftarrow \exp +1023$
FRT $_{12: 63} \leftarrow \mathrm{frac}_{1: 52}$
End
Done

```

\section*{Enabled Exponent Underflow:}

FPSCR \(_{U X} \leftarrow 1\)
sign \(\leftarrow(\text { FRB })_{0}\)
If \((\mathrm{FRB})_{1: 11}=0\) then
Do
        \(\exp \leftarrow-1022\)
        frac \(_{0: 52} \leftarrow 0 \mathrm{bO}\) II \((\mathrm{FRB})_{12: 63}\)
    End
If \((\mathrm{FRB})_{1: 11}>0\) then
    Do
        \(\exp \leftarrow(\mathrm{FRB})_{1: 11}-1023\)
        \(\operatorname{frac}_{0: 52} \leftarrow 0 \mathrm{Ob} 1 \mathrm{II}(\mathrm{FRB})_{12: 63}\)
    End
Normalize operand:
    Do while frac \({ }_{0}=0\)
        \(\exp \leftarrow \exp -1\)
        \(\mathrm{frac}_{0: 52} \leftarrow \mathrm{frac}_{1: 52}\) II Ob0
    End
Round Single(sign,exp,frac \({ }_{0: 52}, 0,0,0\) )
\(\mathrm{FPSCR}_{X X} \leftarrow \mathrm{FPSCR}_{X X} \mid F P S C R_{F I}\)
\(\exp \leftarrow \exp +192\)
\(\mathrm{FRT}_{0} \leftarrow\) sign
FRT \(_{1: 11} \leftarrow \exp +1023\)
\(\mathrm{FRT}_{12: 63} \leftarrow \mathrm{frac}_{1: 52}\)
If sign \(=0\) then FPSCR FPRF \(\leftarrow\) "+ normal number"
If sign = 1 then FPSCR \(_{\text {FPRF }} \leftarrow\) "- normal number"
Done

\section*{Disabled Exponent Overflow:}
\(\mathrm{FPSCR}_{\mathrm{Ox}} \leftarrow 1\)
If FPSCR \({ }_{\text {RN }}=0 \mathrm{~b} 00\) then \(\quad\) Round to Nearest */
Do
If \((F R B)_{0}=0\) then \(\mathrm{FRT} \leftarrow 0 \times 7\) FFO_0000_0000_0000
If \((\text { FRB })_{0}=1\) then FRT \(\leftarrow 0 x F F F 0 \_0000 \_0000 \_0000\)
If \((\mathrm{FRB})_{0}=0\) then \(\mathrm{FPSCR}_{\text {FPRF }} \leftarrow\) "+ infinity"
If \((\mathrm{FRB})_{0}=1\) then FPSCR \(_{\text {FPRF }} \leftarrow{ }^{*}\) - infinity"
End
If \(\mathrm{FPSCR}_{\mathrm{RN}}=0 \mathrm{b01}\) then /* Round toward Zero */
```

        If (FRB)}\mp@subsup{)}{0}{}=0\mathrm{ then FRT }\leftarrow0x47EF_FFFF_E000_0000
        If (FRB)
        If (FRB)}\mp@subsup{)}{0}{}=0\mathrm{ then FPSCR PFPRF }\leftarrow"+\mathrm{ normal number"
        If (FRB)}\mp@subsup{)}{0}{}=1\mathrm{ then FPSCR FPRF }\leftarrow"- normal number"
        End
    If FPSCR 
        Do
            If (FRB)
            If (FRB)}\mp@subsup{)}{0}{}=1\mathrm{ then FRT }\leftarrow0xC7EF_FFFF_E000_0000
            If (FRB)}\mp@subsup{)}{0}{}=0\mathrm{ then FPSCR 
    ```

```

        End
    If FPSCR 
    Do
        If (FRB)
        If (FRB)}\mp@subsup{)}{0}{}=1\mathrm{ then FRT }\leftarrow0xFFFO_0000_0000_0000
        If (FRB)
        If (FRB)
    End
    FPSCR 
    FPSCR 
    FPSCR 
    Done
    Enabled Exponent Overflow:
sign }\leftarrow(FRB\mp@subsup{)}{0}{
exp \leftarrow(FRB) 1:11 - 1023

```

```

    Round Single(sign,exp,frac
    FPSCR 
    Enabled Overflow:
FPSCR
exp \leftarrowexp - 192
FRT
FRT 1:11}\leftarrow\operatorname{exp + 1023
FRT
If sign = 0 then FPSCRRFPRF }\leftarrow "+ normal number"
If sign = 1 then FPSCR FPPRF }\leftarrow\mathrm{ "- normal number"
Done

```

\section*{Zero Operand:}
```

    FRT}\leftarrow(FRB
    ```
    FRT}\leftarrow(FRB
    If (FRB)}\mp@subsup{)}{0}{}=0\mathrm{ then FPSCR 
    If (FRB)}\mp@subsup{)}{0}{=}=1\mathrm{ then FPSCRRPPRF }\leftarrow"- zero"
    FPSCRFR
    Done
Infinity Operand:
    FRT}\leftarrow(FRB
    If (FRB)
    If (FRB)}\mp@subsup{)}{0}{=}=1\mathrm{ then FPSCR 
    FPSCRFR
Done
```


## QNaN Operand:

```
FRT \(\leftarrow(\text { FRB })_{0: 34} \|{ }^{29} 0\)
FPSCR \(_{\text {FPRF }} \leftarrow\) "QNaN"
\(\mathrm{FPSCR}_{\text {FR FI }} \leftarrow 0 \mathrm{~b} 00\)
Done
```


## Version 3.0

```
SNaN Operand:
    FPSCR \(_{\text {VXSNAN }} \leftarrow 1\)
    If \(\mathrm{FPSCR}_{\mathrm{VE}}=0\) then
    Do
        \(\mathrm{FRT}_{0: 11} \leftarrow(\mathrm{FRB})_{0: 11}\)
        \(\mathrm{FRT}_{12} \leftarrow 1\)
        \(\mathrm{FRT}_{13: 63} \leftarrow(\mathrm{FRB})_{13: 34}{ }^{I I}{ }^{29} 0\)
        FPSCR \(_{\text {FPRF }} \leftarrow\) "QNaN"
    End
\(\mathrm{FPSCR}_{\text {FR FI }} \leftarrow 0 \mathrm{~b} 00\)
Done
```


## Normal Operand:

```
sign \(\leftarrow(F R B)_{0}\)
\(\exp \leftarrow(\mathrm{FRB})_{1 \cdot 11}-1023\)
frac \(_{0: 52} \leftarrow 0 \mathrm{~b} 1\) II (FRB) \({ }_{12: 63}\)
Round Single(sign,exp,frac \({ }_{0: 52}, 0,0,0\) )
FPSCR \(_{X X} \leftarrow\) FPSCR \(_{X X} \mid\) FPSCR \({ }_{F I}\)
If exp > 127 and FPSCR \({ }_{\text {OE }}=0\) then go to Disabled Exponent Overflow
If \(\exp >127\) and FPSCR \(\mathrm{OE}=1\) then go to Enabled Overflow
\(\mathrm{FRT}_{0} \leftarrow\) sign
FRT \(_{1: 11} \leftarrow \exp +1023\)
FRT \(_{12: 63} \leftarrow\) frac \(_{1: 52}\)
If sign \(=0\) then FPSCR \(_{\text {FPRF }} \leftarrow\) " + normal number"
If sign \(=1\) then FPSCR FPRRF \(^{\leftarrow}\) "- normal number"
Done
```


## Round Single(sign,exp,frac $\left.{ }_{0: 52}, G, R, X\right)$ :

```
inc \(\leftarrow 0\)
Isb \(\leftarrow \mathrm{frac}_{23}\)
gbit \(\leftarrow \mathrm{frac}_{24}\)
rbit \(\leftarrow \mathrm{frac}_{25}\)
xbit \(\leftarrow\left(\right.\) frac \(\left._{26: 52}\|\mathrm{G}\| R \| \mathrm{X}\right) \neq 0\)
If FPSCR \({ }_{\text {RN }}=0 \mathrm{bOO}\) then Round to Nearest */
Do \(/ *\) comparisons ignore u bits */
If sign || Isb || gbit || rbit || xbit = Obu11uu then inc \(\leftarrow 1\) If sign || Isb || gbit || rbit || xbit = Obu011u then inc \(\leftarrow 1\)
If sign || Isb || gbit || rbit || xbit \(=0\) bu01u1 then inc \(\leftarrow 1\)
End
If FPSCR RN \(=0 \mathrm{~b} 10\) then \(\quad / *\) Round toward + Infinity */
Do /* comparisons ignore u bits */
If sign || Isb || gbit || rbit || xbit \(=0\) bOu1uu then inc \(\leftarrow 1\)
If sign || Isb || gbit I| rbit || xbit \(=0\) bOuu1u then inc \(\leftarrow 1\)
If sign || Isb || gbit || rbit II xbit \(=0 \mathrm{O} 0 \mathrm{uuu} 1\) then inc \(\leftarrow 1\)
End
If \(\mathrm{FPSCR}_{\mathrm{RN}}=0 \mathrm{~b} 11\) then \(\quad / *\) Round toward - Infinity */
Do /* comparisons ignore u bits */
If sign || Isb || gbit || rbit I| xbit = Ob1u1uu then inc \(\leftarrow 1\)
If sign || Isb || gbit || rbit || xbit = 0b1uu1u then inc \(\leftarrow 1\)
If sign || Isb || gbit || rbit || xbit = Ob1uuu1 then inc \(\leftarrow 1\)
End
frac \(_{0: 23} \leftarrow\) frac \(_{0: 23}+\) inc
If carry_out = 1 then
Do
frac \(_{0: 23} \leftarrow 0 \mathrm{~b} 1 \mathrm{II} \mathrm{frac}_{0: 22}\)
\(\exp \leftarrow \exp +1\)
End
frac \(_{24: 52} \leftarrow{ }^{29} 0\)
FPSCR \(_{\text {FR }} \leftarrow\) inc
FPSCR \(_{\text {FI }} \leftarrow\) gbit I rbit | xbit
Return
```


## A. 2 Floating-Point Convert to Integer Model

The following describes algorithmically the operation of the Floating Convert To Integer instructions.

```
if Floating Convert To Integer Word then do
    round_mode }\leftarrow\mathrm{ FPSCR RN
    tgt_precision \leftarrow "32-bit signed integer"
end
if Floating Convert To Integer Word Unsigned then do
    round_mode }\leftarrow\mathrm{ FPSCR RN
    tgt_precision \leftarrow "32-bit unsigned integer"
end
if Floating Convert To Integer Word with round toward Zero then do
    round_mode }\leftarrow00\textrm{b}0
    tgt_precision \leftarrow "32-bit signed integer"
end
if Floating Convert To Integer Word Unsigned with round toward Zero then do
    round_mode }\leftarrow0000
    tgt_precision \leftarrow "32-bit unsigned integer"
end
if Floating Convert To Integer Doubleword then do
    round_mode }\leftarrow\mathrm{ FPSCR RN
    tgt_precision \leftarrow "64-bit signed integer"
end
if Floating Convert To Integer Doubleword Unsigned then do
    round_mode }\leftarrow\mathrm{ FPSCR RN
    tgt_precision \leftarrow "64-bit unsigned integer"
end
if Floating Convert To Integer Doubleword with round toward Zero then do
    round_mode \leftarrow 0b01
    tgt_precision \leftarrow "64-bit signed integer"
end
```

if Floating Convert To Integer Doubleword Unsigned with round toward Zero then do
round_mode $\leftarrow 0 \mathrm{~b} 01$
tgt_precision $\leftarrow$ "64-bit unsigned integer"
end
sign $\leftarrow(F R B)_{0}$
if $(F R B)_{1: 11}=2047$ and $(F R B)_{12: 63}=0$ then goto Infinity Operand
if $(F R B)_{1: 11}=2047$ and $(F R B)_{12}=0$ then goto SNaN Operand
if $(\text { FRB })_{1: 11}=2047$ and $(F R B)_{12}=1$ then goto $Q N a N$ Operand
if $(F R B)_{1: 11}>1086$ then goto Large Operand
if $(\text { FRB })_{1: 11}>0$ then $\exp \leftarrow(\text { FRB })_{1: 11}-1023$ /* exp - bias */
if $(\text { FRB })_{1: 11}=0$ then $\exp \leftarrow-1022$
if $(F R B)_{1: 11}>0$ then frac $_{0: 64} \leftarrow 0 \mathrm{~b} 01| |(F R B)_{12: 63}| | 11_{0}^{11} / *$ normal */
if $(F R B)_{1: 11}=0$ then frac $_{0: 64} \leftarrow 0 \mathrm{~b} 00| |(F R B)_{12: 63}| |{ }^{11} 0$ /* denormal */
gbit || rbit || xbit $\leftarrow 0 b 000$
do i=1,63-exp /* do the loop 0 times if exp = 63 */
frac $_{0}: 64$ || gbit || rbit || xbit $\leftarrow 0$ b0 || frac $0: 64$ || gbit || (rbit | xbit)
end
Round Integer( sign, frac $_{0: 64}$, gbit, rbit, xbit, round_mode )
if sign $=1$ then frac $_{0: 64} \leftarrow \neg$ frac $_{0: 64}+1 / *$ needed leading 0 for $-2^{64}<($ FRB $)<-2^{63} * /$
if tgt_precision = "32-bit signed integer" and frac ${ }_{0: 64}>2^{31}-1$ then goto Large Operand
if tgt_precision $=" 64$-bit signed integer" and $f r a c_{0: 64}>2^{63}-1$ then goto Large Operand
if tgt_precision $=" 32$-bit signed integer" and frac $_{0: 64}<-2^{31}$ then goto Large Operand
if tgt_precision = "64-bit signed integer" and frac ${ }_{0: 64}<-2^{63}$ then goto Large Operand
if tgt_precision $=" 32$-bit unsigned integer" \& frac ${ }_{0: 64}>2^{32}-1$ then goto Large Operand
if tgt_precision = "64-bit unsigned integer" \& frac ${ }_{0: 64}>2^{64}-1$ then goto Large Operand
if tgt_precision = "32-bit unsigned integer" \& frac ${ }_{0: 64}<0$ then goto Large Operand
if tgt_precision = "64-bit unsigned integer" \& frac $0: 64<0$ then goto Large Operand

```
FPSCR
if tgt_precision = "32-bit signed integer" then FRT \leftarrow 0xUUUU_UUUU || frac 33:64
if tgt_precision = "32-bit unsigned integer" then FRT \leftarrow 0xUUUU_UUUU || frac 33:64
if tgt_precision = "64-bit signed integer" then FRT \leftarrow fracce_64
if tgt_precision = "64-bit unsigned integer" then FRT \leftarrow frac}1:6
FPSCR FPRF
done
```

Round Integer( sign, frac ${ }_{0: 64}$, gbit, rbit, xbit, round_mode ):

```
inc \leftarrow 0
if round_mode = 0b00 then do /* Round to Nearest */
    if sign || frac}64 || gbit || rbit || xbit = 0bU11UU then inc \leftarrow < 1
    if sign || frac}64 || gbit || rbit || xbit = 0bU011U then inc \leftarrow < 1
    if sign || frac}64 || gbit || rbit || xbit = 0bU01U1 then inc \leftarrow
end
if round_mode = 0b10 then do /* Round toward +Infinity */
    if sign || frac}64 || gbit || rbit || xbit = 0b0U1UU then inc \leftarrow 1
    if sign || frac}64 || gbit || rbit || xbit = 0b0UU1U then inc \leftarrow < 1
    if sign || frac}64 || gbit || rbit || xbit = 0b0UUU1 then inc \leftarrow 1
end
if round_mode = 0b11 then do /* Round toward -Infinity */
    if sign || frac}64 || gbit || rbit || xbit = 0b1U1UU then inc \leftarrow 1
    if sign || frac}64 || gbit || rbit || xbit = 0b1UU1U then inc \leftarrow < 1
    if sign || frac64 || gbit || rbit || xbit = 0b1UUU1 then inc \leftarrow1
end
frac}00:64 \leftarrow frac 0:64 + in
FPSCR FR }\leftarrow in
FPSCR FI }\leftarrow gbit | rbit | xbi
return
```


## Infinity Operand:

```
FPSCR 
FPSCR FI }\leftarrow < 0b
FPSCR vxcvi }\leftarrow0\textrm{b}
if FPSCRVE = 0 then do
    if tgt_precision = "32-bit signed integer" then do
        if sign=0 then FRT \leftarrow 0xUUUU_UUUU_7FFF_FFFF
        if sign=1 then FRT \leftarrow 0xUUUU_UUUU_8000_0000
    end
    else if tgt_precision = "32-bit unsigned integer" then do
        if sign=0 then FRT \leftarrow 0xUUUU_UUUU_FFFF_FFFF
        if sign=1 then FRT \leftarrow 0xUUUU_UUUU_0000_0000
    end
    else if tgt_precision = "64-bit signed integer" then do
        if sign=0 then FRT \leftarrow 0x7FFF_FFFFF_FFFF_FFFF
        if sign=1 then FRT \leftarrow 0x8000_0000_0000_0000
```

```
    end
    else if tgt_precision = "64-bit unsigned integer" then do
        if sign=0 then FRT \leftarrow0xFFFF_FFFFF_FFFF_FFFF
        if sign=1 then FRT \leftarrow 0x0000_0000_0000_0000
    end
    FPSCR 
end
done
```


## SNaN Operand:

```
    FPSCR 
```

    FPSCR 
    FPSCR 
    FPSCR 
    FPSCRvxsnan \leftarrow 0b1
    FPSCRvxsnan \leftarrow 0b1
    FPSCR VxCVI }\leftarrow00\textrm{b}
    FPSCR VxCVI }\leftarrow00\textrm{b}
    if FPSCR栄 = 0 then do
    if FPSCR栄 = 0 then do
    if tgt_precision = "32-bit signed integer" then FRT \leftarrow 0xUUUU_UUUU_8000_0000
    if tgt_precision = "32-bit signed integer" then FRT \leftarrow 0xUUUU_UUUU_8000_0000
    if tgt_precision = "64-bit signed integer" then FRT \leftarrow 0x8000_0000_0000_0000
    if tgt_precision = "64-bit signed integer" then FRT \leftarrow 0x8000_0000_0000_0000
    if tgt_precision = "32-bit unsigned integer" then FRT \leftarrow 0xUUUU_UUUU_0000_0000
    if tgt_precision = "32-bit unsigned integer" then FRT \leftarrow 0xUUUU_UUUU_0000_0000
    if tgt_precision = "64-bit unsigned integer" then FRT \leftarrow 0x0000_0000_00000_0000
    if tgt_precision = "64-bit unsigned integer" then FRT \leftarrow 0x0000_0000_00000_0000
    FPSCR F
    FPSCR F
    end
end
done

```
done
```


## QNaN Operand:

```
FPSCR FR
```

FPSCR FR
FPSCR
FPSCR
FPSCR
FPSCR
if FPSCRVE = 0 then do
if FPSCRVE = 0 then do
if tgt_precision = "32-bit signed integer" then FRT \leftarrow 0xUUUU_UUUU_8000_0000
if tgt_precision = "32-bit signed integer" then FRT \leftarrow 0xUUUU_UUUU_8000_0000
if tgt_precision = "64-bit signed integer" then FRT \leftarrow 0x8000_0000_0000_0000
if tgt_precision = "64-bit signed integer" then FRT \leftarrow 0x8000_0000_0000_0000
if tgt_precision = "32-bit unsigned integer" then FRT \leftarrow 0xUUUU_UUUU_0000_0000
if tgt_precision = "32-bit unsigned integer" then FRT \leftarrow 0xUUUU_UUUU_0000_0000
if tgt_precision = "64-bit unsigned integer" then FRT \leftarrow 0x0000_0000_0000_0000
if tgt_precision = "64-bit unsigned integer" then FRT \leftarrow 0x0000_0000_0000_0000
FPSCR FPRRF
FPSCR FPRRF
end
end
done
done
Large Operand:
FPSCR
FPSCR
FPSCR
if FPSCR VE = 0 then do
if tgt_precision = "32-bit signed integer" then do
if sign = 0 then FRT \leftarrow 0xUUUU_UUUU_7FFF_FFFF
if sign = 1 then FRT \leftarrow 0xUUUU_UUUU_8000_0000
end
else if tgt_precision = "64-bit signed integer" then do
if sign = 0 then FRT }\leftarrow0x7FFF_FFFF_FFFF_FFF
if sign = 1 then FRT \leftarrow 0x8000_0000_0000_0000
end
else if tgt_precision = "32-bit unsigned integer" then do
if sign = 0 then FRT \leftarrow 0xUUUU_UUUU_FFFF_FFFF
if sign = 1 then FRT \leftarrow 0xUUUU_UUUU_0000_0000
end
else if tgt_precision = "64-bit unsigned integer" then do
if sign = 0 then FRT \leftarrow 0xFFFF_FFFFF_FFFF_FFFF
if sign = 1 then FRT \leftarrow 0x0000_0000_0000_0000
end
FPSCR
end
done

```

\section*{A. 3 Floating-Point Convert from Integer Model}

The following describes algorithmically the operation of the Floating Convert From Integer instructions.
```

    if Floating Convert From Integer Doubleword then do
    tgt_precision \leftarrow "double-precision"
    sign }\leftarrow(FRB)
    exp \leftarrow63
    frac}0:63\leftarrow(FRB
    end
if Floating Convert From Integer Doubleword Single then do
tgt_precision \leftarrow "single-precision"
sign }\leftarrow(FRB)
exp \leftarrow63
frac}0:63\leftarrow(FRB
end
if Floating Convert From Integer Doubleword Unsigned then do
tgt_precision \leftarrow "double-precision"
sign }\leftarrow
exp }\leftarrow6
frac 0:63}\leftarrow(FRB
end
if Floating Convert From Integer Doubleword Unsigned Single then do
tgt_precision \leftarrow "single-precision"
sign }\leftarrow
exp}\leftarrow6
frac}0:63*(FRB
end
if frace:63 = 0 then go to Zero Operand
if sign = 1 then frac 0:63 \leftarrow ᄀfrace:63 + 1
/* do the loop 0 times if (FRB) = max negative 64-bit integer or */
/* if (FRB) = max unsigned 64-bit integer */
do while frace
frac 0:63}\leftarrow\mp@subsup{\textrm{frac}}{1:63 || 0b0}{0
exp \leftarrow exp - 1
end
Round Float( sign, exp, frace:63, RN )
if sign = 0 then FPSCR FPRF }\leftarrow "i+normal number"
if sign = 1 then FPSCR FPRF }\leftarrow"-normal number"
FRT0}\mp@subsup{0}{}{*}\leftarrow\mathrm{ sign
FRT
FRT 12:63}\leftarrow \mp@subsup{\textrm{frac}}{1:52}{
done

```

\section*{Zero Operand:}
```

FPSCR
FPSCR FI }\leftarrow0\textrm{b}0
FPSCR FPRF }\leftarrow"+ zero
FRT \leftarrow 0x0000_0000_0000_0000
done

```
Round Float( sign, exp, frac \(_{0: 63}\), round_mode ):
inc \(\leftarrow 0\)
if tgt_precision = "single-precision" then do
    lsb \(\leftarrow \mathrm{frac}_{23}\)
    gbit \(\leftarrow \mathrm{frac}_{24}\)
    rbit \(\leftarrow \mathrm{frac}_{25}\)
    xbit \(\leftarrow \mathrm{frac}_{26: 63}>0\)
end
else do /* tgt_precision = "double-precision" */
```

    lsb }\leftarrow\mp@subsup{\textrm{frac}}{52}{
    gbit \leftarrow frac}5
    rbit }\leftarrow\mp@subsup{\textrm{frac}}{54}{
    xbit \leftarrow frac 55:63 > 0
    end
if round_mode = 0b00 then do
/* Round to Nearest */
if sign || lsb || gbit || rbit || xbit = 0bU11UU then inc \leftarrow }
if sign || lsb || gbit || rbit || xbit = 0bU011U then inc \leftarrow1
if sign || lsb || gbit || rbit || xbit = 0bU01U1 then inc \leftarrow 1
end
if round_mode = 0b10 then do /* Round toward + Infinity */
if sign || lsb || gbit || rbit || xbit = 0b0U1UU then inc \leftarrow 1
if sign || lsb || gbit || rbit || xbit = 0b0UU1U then inc \leftarrow 1
if sign || lsb || gbit || rbit || xbit = 0b0UUU1 then inc \leftarrow 1
end
if round_mode = 0b11 then do /* Round toward - Infinity */
if sign || lsb || gbit || rbit || xbit = 0b1U1UU then inc \leftarrow
if sign || lsb || gbit || rbit || xbit = 0b1UU1U then inc \leftarrow 1
if sign || lsb || gbit || rbit || xbit = 0b1UUU1 then inc \leftarrow 1
end
if tgt_precision = "single-precision" then
frace:23}\leftarrow\mp@subsup{\textrm{frac}}{0:23}{+ inc
else /* tgt_precision = "double-precision" */
frac}0:52 \leftarrow fraco:52 + in
if carry_out = 1 then exp \leftarrow exp + 1
FPSCR FR
FPSCR
FPSCR
return

```

\section*{A. 4 Floating-Point Round to Integer Model}

The following describes algorithmically the operation of the Floating Round To Integer instructions.
```

If $(F R B)_{1: 11}=2047$ and $(F R B)_{12: 63}=0$, then goto Infinity Operand
If (FRB) $)_{1: 11}=2047$ and $(F R B)_{12}=0$, then goto SNaN Operand
If $(F R B)_{1: 11}=2047$ and $(F R B)_{12}=1$, then goto QNaN Operand
if $(F R B)_{1: 63}=0$ then goto Zero Operand
If (FRB) $1: 11<1023$ then goto Small Operand $/^{*} \exp <0$; Ivaluel < $1^{*} /$
If (FRB) ${ }_{1: 11}>1074$ then goto Large Operand $/{ }^{\star} \exp >51$; integral value */

```
```

sign }\leftarrow(FRB\mp@subsup{)}{0}{
exp \leftarrow(FRB) 1:11-1023 /* exp - bias */
frac}0:52 < Ob1 II (FRB) 12:63
gbit I| rbit |l xbit }\leftarrow0b00

```
Do \(\mathrm{i}=1,52-\exp\)
    \(\operatorname{frac}_{0: 52}\) II gbit || rbit || xbit \(\leftarrow 0\) b0 || frac \(c_{0: 52}\) II gbit || (rbit | xbit)
End
Round Integer (sign, frac \({ }_{0: 52}\), gbit, rbit, xbit)
```

Do $i=2,52-\exp$
frac $_{0: 52} \leftarrow \mathrm{frac}_{1: 52}$ II Ob0
End

```
If \(\mathrm{frac}_{0}=1\), then \(\exp \leftarrow \exp +1\)
Else frac \({ }_{0: 52} \leftarrow\) frac \(_{1: 52}\) II Ob0
\(\mathrm{FRT}_{0} \leftarrow\) sign
FRT \(_{1: 11} \leftarrow \exp +1023\)
FRT \(_{12: 63} \leftarrow\) frac \(_{1: 52}\)
If \((\text { FRT })_{0}=0\) then FPSCR FPRF \(^{\leftarrow} \leftarrow\) " + normal number"
Else FPSCR FPRF \(^{\leftarrow}\) "- normal number"
FPSCR \(_{\text {FR FI }} \leftarrow 0 \mathrm{~b} 00\)
Done
Round Integer(sign, frac0:52, gbit, rbit, xbit):
inc \(\leftarrow 0\)
If inst = Floating Round to Integer Nearest then /* ties away from zero */
    Do /* comparisons ignore u bits */
        If sign || frac \({ }_{52}\) II gbit || rbit I| xbit \(=\) Obuu1uu then inc \(\leftarrow 1\)
        End
If inst = Floating Round to Integer Plus then
        Do /* comparisons ignore u bits */
            If sign II frac \({ }_{52}\) II gbit II rbit II xbit \(=0 b 0\) u1uu then inc \(\leftarrow 1\)
            If sign || frac \({ }_{52}\) II gbit || rbit I| xbit \(=0\) bOuu1u then inc \(\leftarrow 1\)
            If sign || frac 52 || gbit || rbit || xbit \(=0 b 0 u u u 1\) then inc \(\leftarrow 1\)
        End
If inst = Floating Round to Integer Minus then
        Do /* comparisons ignore u bits */
            If sign || frac \({ }_{52}\) I| gbit || rbit || xbit \(=0\) b1u1uu then inc \(\leftarrow 1\)
            If sign I| frac \(\mathrm{c}_{2}\) II gbit || rbit || xbit \(=0\) b1uu1u then inc \(\leftarrow 1\)
            If sign || frac \({ }_{52}\) II gbit || rbit || xbit \(=0\) b1uuu1 then inc \(\leftarrow 1\)
        End
frac \(_{0: 52} \leftarrow\) frac \(_{0: 52}+\) inc
Return

\section*{Infinity Operand:}

FRT \(\leftarrow\) (FRB)
If \((\mathrm{FRB})_{0}=0\) then FPSCR \(_{\text {FPRF }} \leftarrow\) " + infinity"
If \((F R B)_{0}=1\) then FPSCR FPRF \(^{\leftarrow}\) "- infinity"
FPSCR \(_{\text {FR FI }} \leftarrow 0 \mathrm{bOO}\)
Done

\section*{SNaN Operand:}

FPSCR \(_{\text {VXSNAN }} \leftarrow 1\)
If \(F P S C R_{V E}=0\) then
Do
FRT \(\leftarrow\) (FRB)
\(\mathrm{FRT}_{12} \leftarrow 1\)
FPSCR \(_{\text {FPRF }} \leftarrow " Q N a N "\)

\section*{End}

FPSCR \(_{\text {FR FI }} \leftarrow 0\) b00
Done

\section*{QNaN Operand:}

FRT \(\leftarrow\) (FRB)
FPSCR \(_{\text {FPRF }} \leftarrow\) "QNaN"
FPSCR \(_{\text {FR FI }} \leftarrow 0 \mathrm{bOO}\)
Done

\section*{Zero Operand:}

If \((\mathrm{FRB})_{0}=0\) then
Do
FRT \(\leftarrow 0 x 0000 \_0000 \_0000 \_0000\)
FPSCR \(_{\text {FPRF }} \leftarrow\) "+ zero"
End
Else
Do
FRT \(\leftarrow 0 x 8000 \_0000 \_0000 \_0000\)
FPSCR \(_{\text {FPRF }} \leftarrow\) "- zero"
End
FPSCR \(_{\text {FR FI }} \leftarrow 0 \mathrm{~b} 00\)
Done
Small Operand:
If inst = Floating Round to Integer Nearest and (FRB) \()_{1: 11}<1022\) then goto Zero Operand
If inst = Floating Round to Integer Toward Zero then goto Zero Operand
If inst = Floating Round to Integer Plus and (FRB) 0
\(=1\) then goto Zero Operand
If inst = Floating Round to Integer Minus and
\((F R B)_{0}=0\) then goto Zero Operand
If \((\mathrm{FRB})_{0}=0\) then
Do
FRT \(\leftarrow 0 \times 3 F F 0 \_0000 \_0000 \_0000\)
/* value = 1.0 */
FPSCR \(_{\text {FPRF }} \leftarrow\) "+ normal number"
End
Else
Do
FRT \(\leftarrow 0 x B F F 0 \_0000 \_0000 \_0000\)
\(/ *\) value \(=-1.0\) */
FPSCR \(_{\text {FPRF }} \leftarrow\) "- normal number"
End
FPSCR \(_{\text {FRFI }} \leftarrow 0 \mathrm{~b} 00\)
Done

\section*{Large Operand:}

FRT \(\leftarrow\) (FRB)

If \(\mathrm{FRT}_{0}=0\) then \(\mathrm{FPSCR}_{\text {FPRF }} \leftarrow\) "+ normal number"
Else FPSCR FPRF \(^{\leftarrow \text { "- normal number" }}\) FPSCR \(_{\text {FR FI }} \leftarrow 0\) b00
Done

Version 3.0

\title{
Appendix B. Densely Packed Decimal
}

The trailing significand field of the decimal floating-point data format is encoded using Densely Packed Decimal (DPD). DPD encoding is a compression technique which supports the representation of decimal integers of arbitrary length. Translation operates on three Binary Coded Decimal (BCD) digits at a time compressing the 12 bits into 10 bits with an algorithm that
can be applied or reversed using simple Boolean operations. In the following examples, a 3 -digit BCD number is represented as (abcd)(efgh)(ijkm), a 10-bit DPD number is represented as (pqr)(stu)(v)(wxy), and the Boolean operations, \& (AND), I (OR), and \(\neg\) (NOT) are used.

\section*{B. 1 BCD-to-DPD Translation}

The translation from a 3-digit BCD number to a 10-bit DPD can be performed through the following Boolean operations.
```

p = (f \& a \& i \& नe) ( (j \& a \& ᄀi) | (b \& ᄀa)
q = (g\& a \& i \& नe) | (k \& a \& \negi) | (c \& \nega)
r = d
S = (j \& ᄀa \& e \& ᄀi) | (f \& ᄀi \& ᄀe)
(f \& नa \& नe) ( (e \& i)
t = (k \& \nega \& e \& \negi) | (g \& ᄀi \& ᄀe)
(g \& नa \& नe) (a \& i)
u = h
v=a | e | i
w = (\nege \& j \& \negi) | (e \& i) | a
x = (\nega \& k \& ᄀi) (a \& i) e
y = m

```

Alternatively, the following table can be used to perform the translation. The most significant bit of the three BCD digits (left column) is used to select a specific 10 -bit encoding (right column) of the DPD.
\begin{tabular}{|c|cc|}
\hline aei & pqr stu v wxy \\
\hline 000 & bcd fgh 0 jkm \\
\hline 001 & bcd fgh 1 00m \\
\hline 010 & bcd jkh 101 m \\
\hline 011 & bcd 10h 1 11m \\
\hline 100 & jkd fgh 1 10m \\
\hline 101 & fgd 01h 1 11m \\
\hline 110 & jkd 00h 1 11m \\
\hline 111 & \(00 d\) 11h 1 11m \\
\hline
\end{tabular}

The full translation of a 3-digit BCD number (000-999) to a 10-bit DPD is shown in Table 139 on page 793,
with the DPD entries shown in hexadecimal format. The BCD number is produced by replacing ',' in the leftmost column with the corresponding digit along the top row. The table is split into two halves, with the right half being a continuation of the left half.

\section*{B. 2 DPD-to-BCD Translation}

The translation from a 10-bit DPD to a 3-digit BCD number can be performed through the following Boolean operations.
```

a=(\negS \& v \& w) | (t \& v \& w \& S) | (v \& w \& fx)
b}=(\textrm{p}\& S\&\&\&\&\neg)|(p\&\negW)|(p\&\negV
c=(q\& S\& \& \& ᄀt) ( (q\& \&w) ( (q\& \&V)
d = r
e=(v \& ᄀw \& x) | (S \& v \& w \& x)
(\negt \& v \& x \& W)
f = (p \& t \& v \& W \& x \& ~S) | (S \& ᄀx \& v) |
(s \& नV)
g = (q\& \& \& w \& v \& x \& ~S) | (t \& ~x \& v) |
(t \& नV)
h = u
i = (t\& \& \& W \& x) | (S \& V \& W \& X) |
(v \& ᄀW \& \negx)
j=(p\& \&S \& नt \& w \& v) | (S \& v \& नW \& x) |
(p \& w \& ᄀx \& v) (w \& नv)
k=(q\& ~S \& नt \& v \& w) | (t \& v \& नW \& x) |
(q\& v \& W \& नX) | (x \& नV)
m = y

```

Alternatively, the following table can be used to perform the translation. A combination of five bits in the DPD encoding (leftmost column) are used to specify a translation to the 3 -digit BCD encoding. Dashes (-) in the table are don't cares, and can be either one or zero.
\begin{tabular}{|c|c|c|c|}
\hline vwxst & abcd & efgh & ijkm \\
\hline \(0----\) & 0 pqr & 0 stu & \(0 w x y\) \\
\hline \(100--\) & 0 pqr & 0 stu & \(100 y\) \\
\hline \(101--\) & 0 pqr & \(100 u\) & 0 sty \\
\hline \(110--\) & \(100 r\) & 0 stu & \(0 p q y\) \\
\hline 11100 & \(100 r\) & \(100 u\) & \(0 p q y\) \\
\hline 11101 & \(100 r\) & \(0 p q u\) & \(100 y\) \\
\hline 11110 & \(0 p q r\) & \(100 u\) & \(100 y\) \\
\hline 11111 & \(100 r\) & \(100 u\) & \(100 y\) \\
\hline
\end{tabular}

The full translation of the 10-bit DPD to a 3-digit BCD number is shown in Table 140 on page 794. The 10-bit DPD index is produced by concatenating the 6-bit value shown in the left column with the 4-bit index along the top row, both represented in hexadecimal. The values in parentheses are non-preferred translations and are explained further in the following section.

\section*{B. 3 Preferred DPD encoding}

Translating from a 3-digit BCD number (1000 numbers) to a 10-bit DPD encoding (1024 combinations) leaves 24 redundant translations. The 24 redundant combinations are evenly assigned to eight BCD numbers and are shown in the following table, with the non-preferred encoding in parentheses. The preferred encoding is produced by translating a 3 -digit BCD number with the translation table or Boolean operations shown in Section B.1. The redundant DPD encodings are all valid and will be correctly translated to their respective BCD value through the mechanisms provided in Section B.2. For decimal floating-point operations all DPD encodings are recognized as source operands.
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline & \multicolumn{21}{|l|}{Table 139:BCD-to-DPD translation} \\
\hline & 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & & 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\
\hline 00 & 000 & 001 & 002 & 003 & 004 & 005 & 006 & 007 & 008 & 009 & 50 & 280 & 281 & 282 & 283 & 284 & 285 & 286 & 287 & 288 & 289 \\
\hline 01 & 010 & 011 & 012 & 013 & 014 & 015 & 016 & 017 & 018 & 019 & 51 & 290 & 291 & 292 & 293 & 294 & 295 & 296 & 297 & 298 & 299 \\
\hline 02 & 020 & 021 & 022 & 023 & 024 & 025 & 026 & 027 & 028 & 029 & 52 & 2A0 & 2A1 & 2A2 & 2 A 3 & 2A4 & 2A5 & 2A6 & 2A7 & 2A8 & 2A9 \\
\hline 03 & 030 & 031 & 032 & 033 & 034 & 035 & 036 & 037 & 038 & 039 & 53 & 2 BO & 2B1 & 2 B 2 & 2B3 & 2B4 & 2 B 5 & 2B6 & 2B7 & 2B8 & 2B9 \\
\hline 04 & 040 & 041 & 042 & 043 & 044 & 045 & 046 & 047 & 048 & 049 & 54 & 2 CO & 2 C 1 & 2 C 2 & 2 C 3 & 2C4 & 2C5 & 2C6 & 2C7 & 2 C 8 & 2C9 \\
\hline 05 & 050 & 051 & 052 & 053 & 054 & 055 & 56 & 057 & 058 & 059 & 55 & 2D0 & 2D & 2D2 & 2D3 & 2D4 & 5 & 2D & 7 & 8 & 2D9 \\
\hline 06 & 060 & 061 & 062 & 063 & 064 & 065 & 066 & 067 & 068 & 069 & 56 & 2E0 & 2E1 & E2 & 2E3 & 2E4 & 2 E 5 & 2E6 & 2E7 & 2E8 & 2E9 \\
\hline 07 & 070 & 071 & 072 & 073 & 074 & 075 & 076 & 077 & 078 & 079 & 57 & 2F0 & 2F1 & 2 F 2 & 2 F 3 & 2F4 & 2 F 5 & 2F6 & \(2 F 7\) & 2F8 & 2F9 \\
\hline 08 & 00A & 00B & 02A & 02B & 04A & 04B & 06A & 06B & 04E & 04F & 58 & 28A & 28B & 2AA & 2 AB & 2CA & 2CB & 2EA & 2EB & CE & CF \\
\hline 09 & 01A & 01B & 03A & 03B & 05A & 05B & 07A & 07B & 05E & 05F & 59 & 29A & 29B & 2BA & 2BB & 2DA & 2DB & 2FA & 2FB & 2DE & DF \\
\hline 10 & 80 & 081 & 082 & 083 & 084 & 085 & 086 & 087 & 088 & 089 & 60 & 30 & 301 & 302 & 303 & 304 & 305 & 306 & 30 & 308 & 309 \\
\hline 11 & 090 & 091 & 092 & 093 & 094 & 095 & 096 & 097 & 098 & 099 & 61 & 310 & 311 & 31 & 31 & 314 & 315 & 316 & 317 & 318 & 319 \\
\hline 12 & OAO & OA1 & OA2 & OA3 & OA4 & OA5 & 0A6 & OA7 & OA8 & 0A9 & 62 & 320 & 321 & 322 & 323 & 324 & 325 & 326 & 327 & 328 & 329 \\
\hline 13 & OBO & OB1 & OB2 & OB3 & OB4 & 0B5 & 0B6 & OB7 & 0B8 & OB9 & 63 & 330 & 331 & 332 & 333 & 334 & 335 & 336 & 33 & 338 & 339 \\
\hline 14 & OCO & OC1 & OC2 & 0С3 & OC4 & 0C5 & 0C6 & 0C7 & 0C8 & OC9 & 64 & 340 & 341 & 342 & 343 & 344 & 345 & 346 & 347 & 348 & 349 \\
\hline 15 & 0D & OD1 & 0D & OD & OD4 & 0D5 & 0D & 0D7 & 0D8 & OD9 & 65 & 350 & 351 & 352 & 353 & 354 & 355 & 356 & 357 & 358 & 359 \\
\hline 16 & 0 & OE & OE & OE & OE4 & OE & 0E6 & OE & 0E8 & 0E9 & 66 & 360 & 361 & 362 & 363 & 364 & 365 & 366 & 367 & 368 & 369 \\
\hline 17 & OFO & OF1 & OF2 & OF3 & 0F4 & 0F5 & 0F6 & OF7 & 0F8 & 0F9 & 67 & 370 & 371 & 372 & 373 & 374 & 375 & 376 & 377 & 378 & 379 \\
\hline 18 & 08A & 08B & OA & OAB & OCA & OCB & OEA & OEB & OCE & OCF & 68 & 30A & 30B & 32A & 32B & 34A & 34B & 36A & 36B & 34E & 34F \\
\hline 19 & 09A & 09B & OBA & OBB & ODA & ODB & OFA & OFB & ODE & ODF & 69 & 31A & 31B & 33A & 33B & 35A & 35B & 37A & 37B & 35E & 35F \\
\hline 20 & 100 & 101 & 102 & 103 & 104 & 05 & 106 & 107 & 108 & 109 & 70 & 380 & 381 & 382 & 383 & 384 & 385 & 386 & 387 & 388 & 389 \\
\hline 21 & 110 & 111 & 112 & 113 & 114 & 115 & 116 & 117 & 118 & 119 & 71 & 390 & 391 & 392 & 393 & 394 & 395 & 396 & 397 & 398 & 399 \\
\hline 22 & 120 & 121 & 122 & 123 & 124 & 125 & 126 & 127 & 128 & 129 & 72 & 3A0 & 3A1 & 3A2 & 3АЗ & 3A4 & 3A5 & 3A6 & 3A & 3A8 & 3A9 \\
\hline 23 & 130 & 131 & 132 & 133 & 134 & 135 & 136 & 137 & 138 & 139 & 73 & 3B0 & 3B1 & 3B2 & 3B3 & 3B4 & 3B5 & 3B6 & 3B7 & 3B8 & 3B9 \\
\hline 24 & 140 & 141 & 142 & 143 & 144 & 145 & 146 & 147 & 148 & 149 & 74 & 3 CO & 3C1 & 3C2 & 3C3 & 3C4 & 3C5 & 3C6 & 3C7 & 3C8 & 3C9 \\
\hline 25 & 150 & 151 & 152 & 153 & 154 & 155 & 156 & 157 & 158 & 159 & 75 & 3D0 & 3D1 & 3D2 & 3D3 & 3D4 & 3D5 & 3D6 & 3D & 3D8 & 3D9 \\
\hline 26 & 160 & 161 & 162 & 163 & 164 & 65 & 166 & 167 & 168 & 169 & 76 & 3E0 & 3E1 & 3E & E3 & 3E4 & 3E5 & 3E6 & 3E7 & 3E8 & 3E9 \\
\hline 27 & 170 & 171 & 172 & 173 & 174 & 175 & 176 & 177 & 178 & 179 & 77 & 3F0 & 3F1 & 3F2 & 3 F 3 & 3F4 & 3 F 5 & 3F6 & 3F & 3F8 & 3F9 \\
\hline 28 & 10A & 10B & 12A & 12 B & 14A & 14B & 16A & 16B & 14 E & 14F & 78 & 38A & 38B & 3AA & 3 AB & 3CA & 3CB & 3EA & 3E & 3CE & 3C \\
\hline 29 & 11A & 11B & 13A & 13B & 15A & 15B & 17A & 17B & 15E & 15F & 79 & 39A & 39B & 3BA & 3BB & 3DA & 3DB & 3FA & 3FB & 3DE & 3D \\
\hline 30 & 180 & 18 & 182 & 183 & 184 & 185 & 186 & 187 & 88 & 189 & 80 & 000 & 00D & 10 C & 10 & 20 C & 20D & 30C & 30 & 02E & 02F \\
\hline 31 & 190 & 191 & 192 & 193 & 194 & 195 & 196 & 197 & 198 & 199 & 81 & 01C & 01D & 11C & 11 & 21- & 21D & 31C & 31D & 03E & 03F \\
\hline 32 & 1A0 & 1A1 & 1A & 1A3 & 1A & 1A5 & 1A6 & 1A & 1A8 & 1A9 & 82 & 02C & 02D & 12C & 12D & 22 C & 22D & 32C & 32 & 12E & 12F \\
\hline 33 & 1B0 & 1B1 & 1 B 2 & 1B3 & 1B4 & 1B5 & 1B6 & 1B7 & 188 & 1B9 & 83 & 03C & 03D & 13C & 13D & 23C & 23D & 33C & 33D & 13E & 13F \\
\hline 34 & 1C0 & 1C1 & 1 C 2 & 1 C 3 & 1-4 & 1C5 & 1 C 6 & 1C7 & 1 C 8 & 1-9 & 84 & 04C & 04D & 14C & 14D & 24C & 24D & 34C & 34D & 22E & 22F \\
\hline 35 & 1D0 & 1D1 & 1D & 1 D 3 & 1D4 & 1D & 1D6 & 1D7 & 1D8 & 1 D 9 & 85 & 05C & 05D & 15C & 15D & 25C & 25D & 35C & 35D & 23E & F \\
\hline 36 & 1E0 & 1E1 & 1 E 2 & 1E & 1E4 & 1E5 & 1E6 & 1E7 & 1E8 & 1E9 & 86 & 06C & 06D & 16C & 16D & 26C & 26D & 36C & 36D & 32E & 32F \\
\hline 37- & 1F & 1F1 & 1F2 & 1 F3 & 1F4 & 1 F5 & 1F6 & 1 17 & 1F8 & 1 F9 & 87 & 07C & 07D & 17C & 17 & 27C & 27D & 37C & 37D & 33E & 33 F \\
\hline 38 & 18A & 18B & 1AA & 1 AB & 1CA & 1CB & 1EA & 1EB & 1CE & 1CF & 88 & 00E & 00F & 10E & 10 F & 20E & 20F & 30E & 30F & 06E & 06F \\
\hline 39 & 19A & 19B & 1BA & 1BB & 1DA & 1DB & 1FA & 1FB & 1DE & 1DF & 89 & 01E & 01F & 11E & 11F & 21E & 21F & 31E & 31F & 07E & 07F \\
\hline 40 & 200 & 201 & 20 & 203 & 204 & 205 & 20 & 207 & 208 & 209 & 90 & 08C & 08D & 18 & 18 & 28 C & 28D & 38C & 38 & OAE & 0AF \\
\hline 41 & 210 & 211 & 212 & 213 & 214 & 215 & 216 & 217 & 218 & 219 & 91 & 09C & 09D & 19C & 19D & 29C & 29D & 39C & 39D & OBE & OBF \\
\hline 42 & 220 & 221 & 222 & 223 & 224 & 225 & 226 & 227 & 228 & 229 & 92 & OAC & OAD & 1AC & 1AD & 2AC & 2AD & 3AC & 3AD & 1AE & 1AF \\
\hline 43 & 230 & 231 & 23 & 23 & 23 & 235 & 23 & 237 & 238 & 239 & 93 & OBC & OBD & 1BC & 1B & 2BC & 2BD & 3B & 3B & 1BE & 1BF \\
\hline 44 & 240 & 241 & 242 & 243 & 244 & 245 & 246 & 247 & 248 & 249 & 94 & OCC & OCD & 1CC & 1 CD & 2CC & 2CD & 3CC & 3CD & 2AE & 2AF \\
\hline 45 & 250 & 251 & 252 & 253 & 254 & 255 & 256 & 257 & 258 & 259 & 95 & DC & ODD & 1DC & 1DD & 2DC & 2DD & 3DC & 3DD & 2BE & 2BF \\
\hline 46 & 260 & 261 & 262 & 263 & 264 & 265 & 266 & 267 & 268 & 269 & 96 & OEC & OED & 1 EC & 1ED & 2EC & 2ED & 3EC & 3E & 3A & 3AF \\
\hline 47- & 270 & 271 & 272 & 273 & 274 & 275 & 276 & 277 & 278 & 279 & 97 & OFC & OFD & 1FC & 1FD & 2FC & 2FD & 3FC & 3F & 3BE & 3BF \\
\hline 48 & 20A & 20B & 22A & 22B & 24A & 24B & 26A & 26B & 24E & 24F & 98 & 08E & 08F & 18E & 18F & 28E & 28F & 38E & 38 & OEE & OEF \\
\hline 49 & 21A & 21 & 23 & 23 B & 25 & 25B & 27 & 27 & 25E & 25F & 99 & 09 & 09F & 19E & 19 F & 29E & 29 F & 39E & 39 F & OF & OFF \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|c|}
\hline & \multicolumn{16}{|l|}{Table 140: DPD-to-BCD translation} \\
\hline & 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & A & B & C & D & E & F \\
\hline 00 & 000 & 001 & 002 & 003 & 004 & 005 & 006 & 007 & 008 & 009 & 080 & 081 & 800 & 801 & 880 & 881 \\
\hline 01_ & 010 & 011 & 012 & 013 & 014 & 015 & 016 & 017 & 018 & 019 & 090 & 091 & 810 & 811 & 890 & 891 \\
\hline 02 & 020 & 021 & 022 & 023 & 024 & 025 & 026 & 027 & 028 & 029 & 082 & 083 & 820 & 821 & 808 & 809 \\
\hline 03 & 030 & 031 & 032 & 033 & 034 & 035 & 036 & 037 & 038 & 039 & 092 & 093 & 830 & 831 & 818 & 819 \\
\hline 04 & 040 & 041 & 042 & 043 & 044 & 045 & 046 & 047 & 048 & 049 & 084 & 085 & 840 & 841 & 088 & 089 \\
\hline 05 & 050 & 051 & 052 & 053 & 054 & 055 & 056 & 057 & 058 & 059 & 094 & 095 & 850 & 851 & 098 & 099 \\
\hline 06 & 060 & 061 & 062 & 063 & 064 & 065 & 066 & 067 & 068 & 069 & 086 & 087 & 860 & 861 & 888 & 889 \\
\hline 07 & 070 & 071 & 072 & 073 & 074 & 075 & 076 & 077 & 078 & 079 & 096 & 097 & 870 & 871 & 898 & 899 \\
\hline 08 & 100 & 101 & 102 & 103 & 104 & 105 & 106 & 107 & 108 & 109 & 180 & 181 & 900 & 901 & 980 & 981 \\
\hline 09 & 110 & 111 & 112 & 113 & 114 & 115 & 116 & 117 & 118 & 119 & 190 & 191 & 910 & 911 & 990 & 991 \\
\hline OA & 120 & 121 & 122 & 123 & 124 & 125 & 126 & 127 & 128 & 129 & 182 & 183 & 920 & 921 & 908 & 909 \\
\hline OB & 130 & 131 & 132 & 133 & 134 & 135 & 136 & 137 & 138 & 139 & 192 & 193 & 930 & 931 & 918 & 919 \\
\hline OC- & 140 & 141 & 142 & 143 & 144 & 145 & 146 & 147 & 148 & 149 & 184 & 185 & 940 & 941 & 188 & 189 \\
\hline OD_ & 150 & 151 & 152 & 153 & 154 & 155 & 156 & 157 & 158 & 159 & 194 & 195 & 950 & 951 & 198 & 199 \\
\hline OE & 160 & 161 & 162 & 163 & 164 & 165 & 166 & 167 & 168 & 169 & 186 & 187 & 960 & 961 & 988 & 989 \\
\hline 0F & 170 & 171 & 172 & 173 & 174 & 175 & 176 & 177 & 178 & 179 & 196 & 197 & 970 & 971 & 998 & 999 \\
\hline 10 & 200 & 201 & 202 & 203 & 204 & 205 & 206 & 207 & 208 & 209 & 280 & 281 & 802 & 803 & 882 & 883 \\
\hline 11- & 210 & 211 & 212 & 213 & 214 & 215 & 216 & 217 & 218 & 219 & 290 & 291 & 812 & 813 & 892 & 893 \\
\hline 12 & 220 & 221 & 222 & 223 & 224 & 225 & 226 & 227 & 228 & 229 & 282 & 283 & 822 & 823 & 828 & 829 \\
\hline 13 & 230 & 231 & 232 & 233 & 234 & 235 & 236 & 237 & 238 & 239 & 292 & 293 & 832 & 833 & 838 & 839 \\
\hline 14- & 240 & 241 & 242 & 243 & 244 & 245 & 246 & 247 & 248 & 249 & 284 & 285 & 842 & 843 & 288 & 289 \\
\hline 15 & 250 & 251 & 252 & 253 & 254 & 255 & 256 & 257 & 258 & 259 & 294 & 295 & 852 & 853 & 298 & 299 \\
\hline 16 & 260 & 261 & 262 & 263 & 264 & 265 & 266 & 267 & 268 & 269 & 286 & 287 & 862 & 863 & (888) & (889) \\
\hline 17 & 270 & 271 & 272 & 273 & 274 & 275 & 276 & 277 & 278 & 279 & 296 & 297 & 872 & 873 & (898) & (899) \\
\hline 18 & 300 & 301 & 302 & 303 & 304 & 305 & 306 & 307 & 308 & 309 & 380 & 381 & 902 & 903 & 982 & 983 \\
\hline 19 & 310 & 311 & 312 & 313 & 314 & 315 & 316 & 317 & 318 & 319 & 390 & 391 & 912 & 913 & 992 & 993 \\
\hline 1A & 320 & 321 & 322 & 323 & 324 & 325 & 326 & 327 & 328 & 329 & 382 & 383 & 922 & 923 & 928 & 929 \\
\hline 1B & 330 & 331 & 332 & 333 & 334 & 335 & 336 & 337 & 338 & 339 & 392 & 393 & 932 & 933 & 938 & 939 \\
\hline 16- & 340 & 341 & 342 & 343 & 344 & 345 & 346 & 347 & 348 & 349 & 384 & 385 & 942 & 943 & 388 & 389 \\
\hline 1D- & 350 & 351 & 352 & 353 & 354 & 355 & 356 & 357 & 358 & 359 & 394 & 395 & 952 & 953 & 398 & 399 \\
\hline 1E & 360 & 361 & 362 & 363 & 364 & 365 & 366 & 367 & 368 & 369 & 386 & 387 & 962 & 963 & (988) & (989) \\
\hline 1F- & 370 & 371 & 372 & 373 & 374 & 375 & 376 & 377 & 378 & 379 & 396 & 397 & 972 & 973 & (998) & (999) \\
\hline 20- & 400 & 401 & 402 & 403 & 404 & 405 & 406 & 407 & 408 & 409 & 480 & 481 & 804 & 805 & 884 & 885 \\
\hline 21 & 410 & 411 & 412 & 413 & 414 & 415 & 416 & 417 & 418 & 419 & 490 & 491 & 814 & 815 & 894 & 895 \\
\hline 22 & 420 & 421 & 422 & 423 & 424 & 425 & 426 & 427 & 428 & 429 & 482 & 483 & 824 & 825 & 848 & 849 \\
\hline 23 & 430 & 431 & 432 & 433 & 434 & 435 & 436 & 437 & 438 & 439 & 492 & 493 & 834 & 835 & 858 & 859 \\
\hline 24 & 440 & 441 & 442 & 443 & 44 & 445 & 446 & 447 & 448 & 449 & 484 & 485 & 844 & 845 & 488 & 489 \\
\hline 25 & 450 & 451 & 452 & 453 & 454 & 455 & 456 & 457 & 458 & 459 & 494 & 495 & 854 & 855 & 498 & 499 \\
\hline 26 & 460 & 461 & 462 & 463 & 464 & 465 & 466 & 467 & 468 & 469 & 486 & 487 & 864 & 865 & (888) & (889) \\
\hline 27 & 470 & 471 & 472 & 473 & 474 & 475 & 476 & 477 & 478 & 479 & 496 & 497 & 874 & 875 & (898) & (899) \\
\hline 28 & 500 & 501 & 502 & 503 & 504 & 505 & 506 & 507 & 508 & 509 & 580 & 581 & 904 & 905 & 984 & 985 \\
\hline 29 & 510 & 511 & 512 & 513 & 514 & 515 & 516 & 517 & 518 & 519 & 590 & 591 & 914 & 915 & 994 & 995 \\
\hline 2A & 520 & 521 & 522 & 523 & 524 & 525 & 526 & 527 & 528 & 529 & 582 & 583 & 924 & 925 & 948 & 949 \\
\hline 2B & 530 & 531 & 532 & 533 & 534 & 535 & 536 & 537 & 538 & 539 & 592 & 593 & 934 & 935 & 958 & 959 \\
\hline 2C- & 540 & 541 & 542 & 543 & 544 & 545 & 546 & 547 & 548 & 549 & 584 & 585 & 944 & 945 & 588 & 589 \\
\hline 2D & 550 & 551 & 552 & 553 & 554 & 555 & 556 & 557 & 558 & 559 & 594 & 595 & 954 & 955 & 598 & 599 \\
\hline 2E & 560 & 561 & 562 & 563 & 564 & 565 & 566 & 567 & 568 & 569 & 586 & 587 & 964 & 965 & (988) & (989) \\
\hline 2 F & 570 & 571 & 572 & 573 & 574 & 575 & 576 & 577 & 578 & 579 & 596 & 597 & 974 & 975 & (998) & (999) \\
\hline 30- & 600 & 601 & 602 & 603 & 604 & 605 & 606 & 607 & 608 & 609 & 680 & 681 & 806 & 807 & 886 & 887 \\
\hline 31- & 610 & 611 & 612 & 613 & 614 & 615 & 616 & 617 & 618 & 619 & 690 & 691 & 816 & 817 & 896 & 897 \\
\hline 32 & 620 & 621 & 622 & 623 & 624 & 625 & 626 & 627 & 628 & 629 & 682 & 683 & 826 & 827 & 868 & 869 \\
\hline 33 & 630 & 631 & 632 & 633 & 634 & 635 & 636 & 637 & 638 & 639 & 692 & 693 & 836 & 837 & 878 & 879 \\
\hline 34- & 640 & 641 & 642 & 643 & 644 & 645 & 646 & 647 & 648 & 649 & 684 & 685 & 846 & 847 & 688 & 689 \\
\hline 35 & 650 & 651 & 652 & 653 & 654 & 655 & 656 & 657 & 658 & 659 & 694 & 695 & 856 & 857 & 698 & 699 \\
\hline 36 & 660 & 661 & 662 & 663 & 664 & 665 & 666 & 667 & 668 & 669 & 686 & 687 & 866 & 867 & (888) & (889) \\
\hline 37 & 670 & 671 & 672 & 673 & 674 & 675 & 676 & 677 & 678 & 679 & 696 & 697 & 876 & 877 & (898) & (899) \\
\hline 38- & 700 & 701 & 702 & 703 & 704 & 705 & 706 & 707 & 708 & 709 & 780 & 781 & 906 & 907 & 986 & 987 \\
\hline 39- & 710 & 711 & 712 & 713 & 714 & 715 & 716 & 717 & 718 & 719 & 790 & 791 & 916 & 917 & 996 & 997 \\
\hline 3A- & 720 & 721 & 722 & 723 & 724 & 725 & 726 & 727 & 728 & 729 & 782 & 783 & 926 & 927 & 968 & 969 \\
\hline 3B & 730 & 731 & 732 & 733 & 734 & 735 & 736 & 737 & 738 & 739 & 792 & 793 & 936 & 937 & 978 & 979 \\
\hline 3C- & 740 & 741 & 742 & 743 & 744 & 745 & 746 & 747 & 748 & 749 & 784 & 785 & 946 & 947 & 788 & 789 \\
\hline 3D- & 750 & 751 & 752 & 753 & 754 & 755 & 756 & 757 & 758 & 759 & 794 & 795 & 956 & 957 & 798 & 799 \\
\hline 3E- & 760 & 761 & 762 & 763 & 764 & 765 & 766 & 767 & 768 & 769 & 786 & 787 & 966 & 967 & (988) & (989) \\
\hline \(3 \mathrm{~F}_{-}\) & 770 & 771 & 772 & 773 & 774 & 775 & 776 & 777 & 778 & 779 & 796 & 797 & 976 & 977 & (998) & (999) \\
\hline
\end{tabular}

\section*{Appendix C. Assembler Extended Mnemonics}

In order to make assembler language programs simpler to write and easier to understand, a set of extended mnemonics and symbols is provided that defines simple shorthand for the most frequently used forms of Branch Conditional, Compare, Trap, Rotate and Shift, and certain other instructions.

Assemblers should provide the extended mnemonics and symbols listed here, and may provide others.

\section*{C. 1 Symbols}

The following symbols are defined for use in instructions (basic or extended mnemonics) that specify a Condition Register field or a Condition Register bit. The first five (It, ..., un) identify a bit number within a CR field. The remainder (cr0, ..., cr7) identify a CR field. An expression in which a CR field symbol is multiplied by 4 and then added to a bit-number-within-CR-field symbol and 32 can be used to identify a CR bit.
\begin{tabular}{ccl}
\begin{tabular}{c} 
Symbol \\
It
\end{tabular} & \begin{tabular}{c} 
Value
\end{tabular} & \begin{tabular}{l} 
Meaning \\
Less than
\end{tabular} \\
gt & 1 & Greater than \\
eq & 2 & Equal \\
so & 3 & Summary overflow \\
un & 3 & Unordered (after floating-point comparison) \\
cr0 & 0 & CR Field 0 \\
cr1 & 1 & CR Field 1 \\
cr2 & 2 & CR Field 2 \\
cr3 & 3 & CR Field 3 \\
cr4 & 4 & CR Field 4 \\
cr5 & 5 & CR Field 5 \\
cr6 & 6 & CR Field 6 \\
cr7 & 7 & CR Field 7
\end{tabular}

The extended mnemonics in Sections C.2.2 and C. 3 require identification of a CR bit: if one of the CR field symbols is used, it must be multiplied by 4 and added to a bit-number-within-CR-field (value in the range 0-3, explicit or symbolic) and 32. The extended mnemonics in Sections C.2.3 and C. 5 require identification of a CR field: if one of the CR field symbols is used, it must not be multiplied by 4 or added to 32. (For the extended mnemonics in Section C.2.3, the bit number within the CR field is part of the extended mnemonic. The programmer identifies the CR field, and the Assembler does the multiplication and addition required to produce a CR bit number for the BI field of the underlying basic mnemonic.)

\section*{C. 2 Branch Mnemonics}

The mnemonics discussed in this section are variations of the Branch Conditional instructions.
Note: bclr, bclrl, bcctr, and bcctrl each serve as both a basic and an extended mnemonic. The Assembler will recognize a bclr, bclrl, bcctr, or bcctrl mnemonic with three operands as the basic form, and a bclr, bclrl, bcctr, or bcctrl mnemonic with two operands as the extended form. In the extended form the BH operand is omitted and assumed to be Ob00. Similarly, for all the extended mnemonics described in Sections C.2.2-C.2.4 that devolve to any of these four basic mnemonics the BH operand can either be coded or omitted. If it is omitted it is assumed to be Ob00.

\section*{C.2.1 BO and BI Fields}

The 5-bit BO and BI fields control whether the branch is taken. Providing an extended mnemonic for every possible combination of these fields would be neither useful nor practical. The mnemonics described in Sections C.2.2-C.2.4 include the most useful cases. Other cases can be coded using a basic Branch Conditional mnemonic (bc[I] \(\boldsymbol{a}]\), \(\boldsymbol{b c} /[[]], \boldsymbol{b c c t r}[I]\) ) with the appropriate operands.

\section*{C.2.2 Simple Branch Mnemonics}

Instructions using one of the mnemonics in Table 141 that tests a Condition Register bit specify the corresponding bit as the first operand. The symbols defined in Section C. 1 can be used in this operand.

Notice that there are no extended mnemonics for relative and absolute unconditional branches. For these the basic mnemonics b, ba, bl, and bla should be used.
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline Table 141:Simple branch mne & nics & & & & & & & \\
\hline \multirow[b]{2}{*}{Branch Semantics} & \multicolumn{4}{|c|}{LR not Set} & \multicolumn{4}{|c|}{LR Set} \\
\hline & bc Relative & \begin{tabular}{l}
bca \\
Absolute
\end{tabular} & \[
\begin{gathered}
\text { bclr } \\
\text { To LR }
\end{gathered}
\] & \begin{tabular}{l}
bcctr \\
To CTR
\end{tabular} & \begin{tabular}{l}
bcl \\
Relative
\end{tabular} & \begin{tabular}{l}
bcla \\
Absolute
\end{tabular} & \[
\begin{aligned}
& \text { bcIrl } \\
& \text { To LR }
\end{aligned}
\] & \[
\begin{aligned}
& \text { bcctrl } \\
& \text { To CTR }
\end{aligned}
\] \\
\hline Branch unconditionally & - & - & blr & bctr & - & - & blrl & bctrl \\
\hline Branch if \(\mathrm{CR}_{\mathrm{Bl}}=1\) & bt & bta & btlr & btctr & btl & btla & bt|rl & btctrl \\
\hline Branch if \(\mathrm{CR}_{\mathrm{BI}}=0\) & bf & bfa & bflr & bfctr & bfl & bfla & bflrl & bfctrl \\
\hline Decrement CTR, branch if CTR nonzero & bdnz & bdnza & bdnzlr & - & bdnzl & bdnzla & bdnzıIr & - \\
\hline Decrement CTR, branch if CTR nonzero and \(\mathrm{CR}_{\mathrm{BI}}=1\) & bdnzt & bdnzta & bdnztlr & - & bdnztl & bdnztla & bdnzt|rl & - \\
\hline Decrement CTR, branch if CTR nonzero and \(\mathrm{CR}_{\mathrm{BI}}=0\) & bdnzf & bdnzfa & bdnzflr & - & bdnzfl & bdnzfla & bdnzflıl & - \\
\hline \begin{tabular}{|l|}
\hline \begin{tabular}{l} 
Decrement CTR, branch if \\
CTR zero
\end{tabular} \\
\hline
\end{tabular} & bdz & bdza & bdzlr & - & bdzl & bdzla & bdzırı & - \\
\hline Decrement CTR, branch if CTR zero and \(\mathrm{CR}_{\mathrm{BI}}=1\) & bdzt & bdzta & bdztlr & - & bdztl & bdztla & bdzt|rl & - \\
\hline Decrement CTR, branch if CTR zero and \(\mathrm{CR}_{\mathrm{BI}}=0\) & bdzf & bdzfa & bdzflr & - & bdzfl & bdzfla & bdzflrl & - \\
\hline
\end{tabular}

\section*{Examples}
1. Decrement CTR and branch if it is still nonzero (closure of a loop controlled by a count loaded into CTR).
bdnz target (equivalent to: bc 16,0,target)
2. Same as (1) but branch only if CTR is nonzero and condition in CRO is "equal".
\[
\text { bdnzt eq,target } \quad \text { (equivalent to: bc } 8,2 \text {,target) }
\]
3. Same as (2), but "equal" condition is in CR5.
bdnzt \(4 \times\) cr5+eq,target (equivalent to: bc 8,22 ,target)
4. Branch if bit 59 of \(C R\) is 0 .
bf 27,target (equivalent to: bc 4,27,target)
5. Same as (4), but set the Link Register. This is a form of conditional "call".
bfl
27,target (equivalent to:
bcl 4,27,target)

\section*{C.2.3 Branch Mnemonics Incorporating Conditions}

In the mnemonics defined in Table 142, the test of a bit in a Condition Register field is encoded in the mnemonic.
Instructions using the mnemonics in Table 142 specify the CR field as an optional first operand. One of the CR field symbols defined in Section C. 1 can be used for this operand. If the CR field being tested is CR Field 0, this operand need not be specified unless the resulting basic mnemonic is bclr[I] or bcctr[I] and the BH operand is specified.
A standard set of codes has been adopted for the most common combinations of branch conditions.
\begin{tabular}{ll} 
Code & Meaning \\
It & Less than \\
le & Less than or equal \\
eq & Equal \\
ge & Greater than or equal \\
gt & Greater than \\
nl & Not less than \\
ne & Not equal \\
ng & Not greater than \\
so & Summary overflow \\
ns & Not summary overflow \\
un & Unordered (after floating-point comparison) \\
nu & Not unordered (after floating-point comparison)
\end{tabular}

These codes are reflected in the mnemonics shown in Table 142.
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \multirow[b]{2}{*}{Branch Semantics} & \multicolumn{4}{|c|}{LR not Set} & \multicolumn{4}{|c|}{LR Set} \\
\hline & bc Relative & bca Absolute & \[
\begin{gathered}
\text { bclr } \\
\text { To LR }
\end{gathered}
\] & \[
\begin{array}{|c|}
\hline \text { bcctr } \\
\text { To CTR }
\end{array}
\] & bcl Relative & bcla Absolute & \[
\begin{aligned}
& \text { bcIrI } \\
& \text { To LR }
\end{aligned}
\] & bcctrl To CTR \\
\hline Branch if less than & blt & blta & blttr & bltctr & bltt & blta & bltirl & bltctrl \\
\hline Branch if less than or equal & ble & blea & blelr & blectr & blel & blela & blelrl & blectrl \\
\hline Branch if equal & beq & beqa & beqlr & beqctr & beql & beqla & beqlrl & beqctrl \\
\hline Branch if greater than or equal & bge & bgea & bgelr & bgectr & bgel & bgela & bgelrl & bgectrl \\
\hline Branch if greater than & bgt & bgta & bgtlr & bgtctr & bgtl & bgtla & bgtlrl & bgtctrl \\
\hline Branch if not less than & bnl & bnla & bnllr & bnlctr & bnll & bnlla & bnllrl & bnlctrl \\
\hline Branch if not equal & bne & bnea & bnelr & bnectr & bnel & bnela & bnelrl & bnectrl \\
\hline Branch if not greater than & bng & bnga & bnglr & bngctr & bngl & bngla & bnglrl & bngctrl \\
\hline Branch if summary overflow & bso & bsoa & bsolr & bsoctr & bsol & bsola & bsolrl & bsoctrl \\
\hline Branch if not summary overflow & bns & bnsa & bnslr & bnsctr & bnsl & bnsla & bnsirl & bnsctrl \\
\hline Branch if unordered & bun & buna & bunlr & bunctr & bunl & bunla & bunirl & bunctrl \\
\hline Branch if not unordered & bnu & bnua & bnulr & bnuctr & bnul & bnula & bnulrl & bnuctrl \\
\hline
\end{tabular}

\section*{Examples}
1. Branch if CRO reflects condition "not equal".
bne target (equivalent to: bc 4,2,target)
2. Same as (1), but condition is in CR3.

\section*{Version 3.0}
bne cr3,target (equivalent to: bc 4,14,target)
3. Branch to an absolute target if CR4 specifies "greater than", setting the Link Register. This is a form of conditional "call".
bgtla cr4,target (equivalent to: bcla 12,17,target)
4. Same as (3), but target address is in the Count Register.
bgtctrl cr4 (equivalent to: bcctrl \(12,17,0\) )

\section*{C.2.4 Branch Prediction}

Software can use the "at" bits of Branch Conditional instructions to provide a hint to the processor about the behavior of the branch. If, for a given such instruction, the branch is almost always taken or almost always not taken, a suffix can be added to the mnemonic indicating the value to be used for the "at" bits.
+ Predict branch to be taken (at=0b11)
- Predict branch not to be taken (at=0b10)

Such a suffix can be added to any Branch Conditional mnemonic, either basic or extended, that tests either the Count Register or a CR bit (but not both). Assemblers should use \(0 b 00\) as the default value for the "at" bits, indicating that software has offered no prediction.

\section*{Examples}
1. Branch if CRO reflects condition "less than", specifying that the branch should be predicted to be taken.
blt+ target
2. Same as (1), but target address is in the Link Register and the branch should be predicted not to be taken.
blttr-

\section*{C. 3 Condition Register Logical Mnemonics}

The Condition Register Logical instructions can be used to set (to 1), clear (to 0), copy, or invert a given Condition Register bit. Extended mnemonics are provided that allow these operations to be coded easily.

Table 143:Condition Register logical mnemonics
\begin{tabular}{|l|l|l|}
\hline Operation & Extended Mnemonic & Equivalent to \\
\hline Condition Register set & crset bx & creqv bx,bx,bx \\
\hline Condition Register clear & crclr bx & crxor bx,bx,bx \\
\hline Condition Register move & crmove bx,by & cror bx,by,by \\
\hline Condition Register not & crnot bx,by & crnor bx,by,by \\
\hline
\end{tabular}

The symbols defined in Section C. 1 can be used to identify the Condition Register bits.

\section*{Examples}
1. Set CR bit 57.
crset 25
(equivalent to: creqv \(25,25,25\) )
2. Clear the SO bit of CR 0 . crclr so (equivalent to: crxor \(3,3,3\) )
3. Same as (2), but SO bit to be cleared is in CR3.
crclr \(4 \times\) cr3+so (equivalent to: crxor \(15,15,15\) )
4. Invert the EQ bit.
crnot eq,eq (equivalent to: crnor 2,2,2)
5. Same as (4), but EQ bit to be inverted is in CR4, and the result is to be placed into the EQ bit of CR5.
crnot \(4 \times\) cr5+eq, \(4 \times\) cr \(4+e q \quad\) (equivalent to: crnor \(22,18,18\) )

\section*{C. 4 Subtract Mnemonics}

\section*{C.4.1 Subtract Immediate}

Although there is no "Subtract Immediate" instruction, its effect can be achieved by using an Add Immediate instruction with the immediate operand negated. Extended mnemonics are provided that include this negation, making the intent of the computation clearer.
\begin{tabular}{lllll} 
subi & \(R x, R y\), value & (equivalent to: & addi & \(R x, R y\), -value) \\
subis & \(R x, R y\), value & (equivalent to: & addis & \(R x, R y\), -value) \\
subic & \(R x, R y\), value & (equivalent to: & addic & \(R x, R y\), -value) \\
subic. & \(R x, R y\), value & (equivalent to: & addic. & \(R x, R y\), -value)
\end{tabular}

\section*{C.4.2 Subtract}

The Subtract From instructions subtract the second operand (RA) from the third (RB). Extended mnemonics are provided that use the more "normal" order, in which the third operand is subtracted from the second. Both these mnemonics can be coded with a final "o" and/or "." to cause the OE and/or Rc bit to be set in the underlying instruction.
\begin{tabular}{lllll} 
sub & \(R x, R y, R z\) & (equivalent to: & subf & \(R x, R z, R y\) ) \\
subc & \(R x, R y, R z\) & (equivalent to: & subfc & \(R x, R z, R y\) )
\end{tabular}

\section*{Version 3.0}

\section*{C. 5 Compare Mnemonics}

The \(L\) field in the fixed-point Compare instructions controls whether the operands are treated as 64-bit quantities or as 32 -bit quantities. Extended mnemonics are provided that represent the \(L\) value in the mnemonic rather than requiring it to be coded as a numeric operand.

The BF field can be omitted if the result of the comparison is to be placed into CR Field 0 . Otherwise the target CR field must be specified as the first operand. One of the CR field symbols defined in Section C. 1 can be used for this operand.

I Note: The Assembler will recognize a basic Compare mnemonic with three operands, and will generate the instruction with \(L=0\). Thus the Assembler must require that the BF field, which normally can be omitted when CR Field 0 is the target, be specified explicitly if \(L\) is.

\section*{C.5.1 Doubleword Comparisons}
\begin{tabular}{|l|l|l|}
\hline \multicolumn{3}{|l|}{ Table 144:Doubleword compare mnemonics } \\
\hline Operation & Extended Mnemonic & Equivalent to \\
\hline Compare doubleword immediate & cmpdi bf,ra,si & cmpi bf,1,ra,si \\
\hline Compare doubleword & cmpd bf,ra,rb & cmp bf, \(1, \mathrm{ra}, \mathrm{rb}\) \\
\hline Compare logical doubleword immediate & cmpldi bf,ra,ui & \(\mathrm{cmpli} \mathrm{bf}, 1, \mathrm{ra}, \mathrm{ui}\) \\
\hline Compare logical doubleword & cmpld bf,ra,rb & cmpl bf,1,ra,rb \\
\hline
\end{tabular}

\section*{Examples}
1. Compare register \(R x\) and immediate value 100 as unsigned 64 -bit integers and place result into CRO.
cmpldi
\(R x, 100\)
(equivalent to: cmpli
\(0,1, R x, 100)\)
2. Same as (1), but place result into CR4.
cmpldi cr4,Rx,100 (equivalent to: cmpli 4,1,Rx,100)
3. Compare registers Rx and Ry as signed 64-bit integers and place result into CRO.
cmpd
Rx,Ry
(equivalent to: cmp
0,1,Rx,Ry)

\section*{C.5.2 Word Comparisons}

Table 145:Word compare mnemonics
\begin{tabular}{|l|l|l|}
\hline Operation & Extended Mnemonic & Equivalent to \\
\hline Compare word immediate & cmpwi bf,ra,si & cmpi bf,0,ra,si \\
\hline Compare word & cmpw bf,ra,rb & cmp bf,0,ra,rb \\
\hline Compare logical word immediate & cmplwi bf,ra,ui & cmpli bf,0,ra,ui \\
\hline Compare logical word & cmplw bf,ra,rb & cmpl bf,0,ra,rb \\
\hline
\end{tabular}

\section*{Examples}
1. Compare bits \(32: 63\) of register \(R x\) and immediate value 100 as signed 32 -bit integers and place result into CR0.
cmpwi \(\mathrm{Rx}, 100\) (equivalent to: cmpi 0,0,Rx,100)
2. Same as (1), but place result into CR4.
cmpwi cr4,Rx,100 (equivalent to: cmpi 4,0,Rx,100)
3. Compare bits 32:63 of registers Rx and Ry as unsigned 32-bit integers and place result into CRO.
cmplw Rx,Ry (equivalent to: \(\mathrm{cmpl} \quad 0,0, R x, R y\) )

\section*{Version 3.0}

\section*{C. 6 Trap Mnemonics}

The mnemonics defined in Table 146 are variations of the Trap instructions, with the most useful values of TO represented in the mnemonic rather than specified as a numeric operand.

A standard set of codes has been adopted for the most common combinations of trap conditions.
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline Code & Meaning & TO encoding & < & > & & & > \\
\hline It & Less than & 16 & 1 & 0 & 0 & 0 & 0 \\
\hline le & Less than or equal & 20 & 1 & 0 & 1 & 0 & 0 \\
\hline eq & Equal & 4 & 0 & 0 & 1 & 0 & 0 \\
\hline ge & Greater than or equal & 12 & 0 & 1 & 1 & 0 & 0 \\
\hline gt & Greater than & 8 & 0 & 1 & 0 & 0 & 0 \\
\hline nl & Not less than & 12 & 0 & 1 & 1 & 0 & 0 \\
\hline ne & Not equal & 24 & 1 & 1 & 0 & 0 & 0 \\
\hline ng & Not greater than & 20 & 1 & 0 & 1 & 0 & 0 \\
\hline lit & Logically less than & 2 & 0 & 0 & 0 & 1 & 0 \\
\hline 1 le & Logically less than or equal & 6 & 0 & 0 & 1 & 1 & 0 \\
\hline Ige & Logically greater than or equal & 5 & 0 & 0 & 1 & 0 & 1 \\
\hline Igt & Logically greater than & 1 & 0 & 0 & 0 & 0 & \\
\hline Inl & Logically not less than & 5 & 0 & 0 & 1 & 0 & 1 \\
\hline Ing & Logically not greater than & 6 & 0 & & 1 & 1 & 0 \\
\hline u & Unconditionally with parameters & 31 & 1 & 1 & 1 & 1 & 1 \\
\hline (none) & Unconditional & 31 & 1 & 1 & 1 & 1 & 1 \\
\hline
\end{tabular}

These codes are reflected in the mnemonics shown in Table 146.
\begin{tabular}{|c|c|c|c|c|}
\hline \multicolumn{5}{|l|}{Table 146:Trap mnemonics} \\
\hline \multirow[b]{2}{*}{Trap Semantics} & \multicolumn{2}{|l|}{64-bit Comparison} & \multicolumn{2}{|l|}{32-bit Comparison} \\
\hline & \(t d i\) Immediate & \(t d\) Register & twi Immediate & \(t w\) Register \\
\hline Trap unconditionally & - & - & - & trap \\
\hline Trap unconditionally with parameters & tdui & tdu & twui & twu \\
\hline Trap if less than & tdllti & tdlt & twlti & twit \\
\hline Trap if less than or equal & tdlei & tdle & twlei & twle \\
\hline Trap if equal & tdeqi & tdeq & tweqi & tweq \\
\hline Trap if greater than or equal & tdgei & tdge & twgei & twge \\
\hline Trap if greater than & tdgti & tdgt & twgti & twgt \\
\hline Trap if not less than & tdnli & tdnl & twnli & twnl \\
\hline Trap if not equal & tdnei & tdne & twnei & twne \\
\hline Trap if not greater than & tdngi & tdng & twngi & twng \\
\hline Trap if logically less than & tdllti & tdillt & twllti & twllt \\
\hline Trap if logically less than or equal & tdllei & tdlle & twllei & twlle \\
\hline Trap if logically greater than or equal & tdlgei & tdlge & twigei & twlge \\
\hline Trap if logically greater than & tdlgti & tdlgt & twlgti & twigt \\
\hline Trap if logically not less than & tdInli & tdlnı & twinli & twinl \\
\hline Trap if logically not greater than & tdlngi & tdling & twlngi & twing \\
\hline
\end{tabular}

\section*{Version 3.0}

\section*{Examples}
1. Trap if register \(R x\) is not 0 .
tdnei Rx,0 (equivalent to: tdi 24,Rx,0)
2. Same as (1), but comparison is to register Ry.
tdne \(R x, R y \quad\) (equivalent to: td 24,Rx,Ry)
3. Trap if bits \(32: 63\) of register \(R x\), considered as a 32 -bit quantity, are logically greater than \(0 x 7 F F\).
twigti \(R x, 0 x 7 F F \quad\) (equivalent to: twi \(1, R x, 0 x 7 F F\) )
4. Trap unconditionally.
trap (equivalent to: tw 31,0,0)
5. Trap unconditionally with immediate parameters \(R x\) and Ry
tdu Rx,Ry (equivalent to: td 31,Rx,Ry)

\section*{C. 7 Integer Select Mnemonics}

The mnemonics defined in Table 147, "Integer Select mnemonics," on page 802 are variations of the Integer Select instructions, with the most useful values of BC represented in the mnemonic rather than specified as a numeric operand..
\begin{tabular}{ll} 
Code & Meaning \\
It & Less than \\
eq & Equal \\
gt & Greater than
\end{tabular}

These codes are reflected in the mnemonics shown in Table 147.
\begin{tabular}{|l|c|}
\hline Table 147: Integer Select mnemonics & \\
\hline \multicolumn{1}{|c|}{ Select semantics } & \begin{tabular}{c} 
isel \\
extended \\
mnemonic
\end{tabular} \\
\hline Integer Select if less than & isellt \\
\hline Integer Select if equal & iseleq \\
\hline Integer Select if greater than & iselgt \\
\hline
\end{tabular}

\section*{Examples}
1. Set register Rx to Ry if the LT bit is set in CRO, and to Rz otherwise.
isellt \(R x, R y, R z \quad\) (equivalent to: isel \(R x, R y, R z, 0\) )
2. Set register Rx to Ry if the GT bit is set in CR0, and to Rz otherwise.
iselgt \(R x, R y, R z \quad\) (equivalent to: isel \(R x, R y, R z, 1)\)
3. Set register \(R x\) to \(R y\) if the EQ bit is set in CRO, and to Rz otherwise.
iseleq \(R x, R y, R z \quad\) (equivalent to: isel \(R x, R y, R z, 2\) )

\section*{C. 8 Rotate and Shift Mnemonics}

The Rotate and Shift instructions provide powerful and general ways to manipulate register contents, but can be difficult to understand. Extended mnemonics are provided that allow some of the simpler operations to be coded easily.
Mnemonics are provided for the following types of operation.
Extract Select a field of \(n\) bits starting at bit position \(b\) in the source register; left or right justify this field in the target register; clear all other bits of the target register to 0 .

Insert Select a left-justified or right-justified field of \(n\) bits in the source register; insert this field starting at bit position b of the target register; leave other bits of the target register unchanged. (No extended mnemonic is provided for insertion of a left-justified field when operating on doublewords, because such an insertion requires more than one instruction.)
Rotate Rotate the contents of a register right or left n bits without masking.
Shift Shift the contents of a register right or left n bits, clearing vacated bits to 0 (logical shift).
Clear Clear the leftmost or rightmost n bits of a register to 0 .

\section*{Clear left and shift left}

Clear the leftmost \(b\) bits of a register, then shift the register left by \(n\) bits. This operation can be used to scale a (known nonnegative) array index by the width of an element.

\section*{C.8.1 Operations on Doublewords}

All these mnemonics can be coded with a final "." to cause the Rc bit to be set in the underlying instruction.
\begin{tabular}{|c|c|c|}
\hline \multicolumn{3}{|l|}{Table 148:Doubleword rotate and shift mnemonics} \\
\hline Operation & Extended Mnemonic & Equivalent to \\
\hline Extract and left justify immediate & extldi ra,rs,n,b ( \(\mathrm{n}>0\) ) & rldicr ra,rs, b, \(\mathrm{n}-1\) \\
\hline Extract and right justify immediate & extrdi ra,rs,n,b ( \(n>0\) ) & rldicl ra,rs, \(\mathrm{b}+\mathrm{n}, 64-\mathrm{n}\) \\
\hline Insert from right immediate & insrdi ra,rs,n,b ( \(\mathrm{n}>0\) ) & rldimi ra,rs,64-(b+n), b \\
\hline Rotate left immediate & rotldi ra,rs,n & rldicl ra,rs,n,0 \\
\hline Rotate right immediate & rotrdi ra,rs,n & rldicl ra,rs,64-n,0 \\
\hline Rotate left & rotld ra,rs,rb & rldcl ra,rs,rb,0 \\
\hline Shift left immediate & sldi ra,rs,n ( n < 64) & rldicr ra,rs,n,63-n \\
\hline Shift right immediate & srdi ra,rs, n ( \(\mathrm{n}<64\) ) & rldicl ra,rs,64-n,n \\
\hline Clear left immediate & Clrldi ra,rs, n ( \(\mathrm{n}<64\) ) & rldicl ra,rs, \(0, \mathrm{n}\) \\
\hline Clear right immediate & clrrdi ra,rs, n ( \(\mathrm{n}<64\) ) & rldicr ra,rs,0,63-n \\
\hline Clear left and shift left immediate & clrlsldi ra,rs,b,n ( \(\mathrm{n}<=\mathrm{b}<64\) ) & rldic ra,rs,n,b-n \\
\hline
\end{tabular}

\section*{Examples}
1. Extract the sign bit (bit 0) of register Ry and place the result right-justified into register Rx.
\[
\text { extrdi } \quad R x, R y, 1,0 \quad \text { (equivalent to: } \quad \text { rldicl } \quad R x, R y, 1,63)
\]
2. Insert the bit extracted in (1) into the sign bit (bit 0 ) of register Rz.
insrdi \(R z, R x, 1,0 \quad\) (equivalent to: rldimi \(R z, R x, 63,0\) )
3. Shift the contents of register Rx left 8 bits.
sldi
\(R x, R x, 8\)
(equivalent to: rldicr
\(R x, R x, 8,55)\)
4. Clear the high-order 32 bits of register Ry and place the result into register Rx.
clrldi Rx,Ry,32 (equivalent to: rldicl Rx,Ry,0,32)

\section*{Version 3.0}

\section*{C.8.2 Operations on Words}

All these mnemonics can be coded with a final "." to cause the Rc bit to be set in the underlying instruction. The operations as described above apply to the low-order 32 bits of the registers, as if the registers were 32-bit registers. The Insert operations either preserve the high-order 32 bits of the target register or place rotated data there; the other operations clear these bits.

Table 149:Word rotate and shift mnemonics
\begin{tabular}{|c|c|c|}
\hline Operation & Extended Mnemonic & Equivalent to \\
\hline Extract and left justify immediate & extlwi ra,rs, \(\mathrm{n}, \mathrm{b} \quad(\mathrm{n}>0)\) & rlwinm ra,rs,b,0,n-1 \\
\hline Extract and right justify immediate & extrwi ra,rs,n,b ( \(\mathrm{n}>0\) ) & rlwinm ra,rs,b+n,32-n,31 \\
\hline Insert from left immediate & inslwi ra,rs,n,b \(\quad(\mathrm{n}>0)\) & rlwimi ra,rs,32-b,b, (b+n)-1 \\
\hline Insert from right immediate & insrwi ra,rs,n,b \(\quad(\mathrm{n}>0)\) & rlwimi ra,rs, \(32-(\mathrm{b}+\mathrm{n}\) ), b, (b+n)-1 \\
\hline Rotate left immediate & rotlwi ra,rs,n & rlwinm ra,rs,n,0,31 \\
\hline Rotate right immediate & rotrwi ra,rs,n & rlwinm ra,rs,32-n,0,31 \\
\hline Rotate left & rotlw ra,rs,rb & rlwnm ra,rs,rb,0,31 \\
\hline Shift left immediate & slwi ra,rs, \(\mathrm{n} \quad(\mathrm{n}<32)\) & rlwinm ra,rs,n,0,31-n \\
\hline Shift right immediate & srwi ra,rs,n \(\quad(\mathrm{n}<32)\) & rlwinm ra,rs,32-n,n,31 \\
\hline Clear left immediate & clrlwi ra,rs, n ( \(\mathrm{n}<32)\) & rlwinm ra,rs,0,n,31 \\
\hline Clear right immediate & clrrwi ra,rs, n ( \(\mathrm{n}<32)\) & rlwinm ra,rs, 0,0,31-n \\
\hline Clear left and shift left immediate & clrlslwi \(\quad\) ra,rs,b, \(\mathrm{n} \quad(\mathrm{n} \leq \mathrm{b}<32)\) & rlwinm ra,rs,n,b-n,31-n \\
\hline
\end{tabular}

\section*{Examples}
1. Extract the sign bit (bit 32) of register Ry and place the result right-justified into register Rx.
\[
\text { extrwi } R x, R y, 1,0 \quad \text { (equivalent to: } \quad \text { rlwinm } \quad R x, R y, 1,31,31 \text { ) }
\]
2. Insert the bit extracted in (1) into the sign bit (bit 32) of register Rz.
insrwi \(R z, R x, 1,0 \quad\) (equivalent to: rlwimi \(R z, R x, 31,0,0)\)
3. Shift the contents of register Rx left 8 bits, clearing the high-order 32 bits.
slwi \(R x, R x, 8 \quad\) (equivalent to: rlwinm \(R x, R x, 8,0,23\) )
4. Clear the high-order 16 bits of the low-order 32 bits of register Ry and place the result into register Rx, clearing the high-order 32 bits of register Rx.

Clrlwi \(\quad R x, R y, 16 \quad\) (equivalent to: rlwinm \(R x, R y, 0,16,31\) )

\section*{C. 9 Move To/From Special Purpose Register Mnemonics}

The \(\boldsymbol{m} t \boldsymbol{s p r}\) and \(\boldsymbol{m f s p r}\) instructions specify a Special Purpose Register (SPR) as a numeric operand. Extended mnemonics are provided that represent the SPR in the mnemonic rather than requiring it to be coded as an operand.

Table 150:Extended mnemonics for moving to/from an SPR
\begin{tabular}{|c|c|c|c|c|}
\hline \multirow{2}{*}{Special Purpose Register} & \multicolumn{2}{|c|}{Move To SPR} & \multicolumn{2}{|r|}{Move From SPR} \\
\hline & Extended & Equivalent to & Extended & Equivalent to \\
\hline XER & mtxer Rx & mtspr 1,Rx & mfxer Rx & mfspr Rx, 1 \\
\hline DSCR & mtudscr Rx & mtspr 3,Rx & mfudscr Rx & mfspr Rx,3 \\
\hline LR & mtlr Rx & mtspr 8,Rx & mflr Rx & mfspr Rx, 8 \\
\hline CTR & mtctr Rx & mtspr 9,Rx & mfctr Rx & mfspr Rx, 9 \\
\hline AMR & mtuamr Rx & mtspr 13,Rx & mfuamr Rx & mfspr Rx, 13 \\
\hline TFHAR & mttfhar Rx & mtspr 128,Rx & mftfhar Rx & mfspr Rx, 128 \\
\hline TFIAR & mtffiar Rx & mtspr 129,Rx & mftfiar Rx & mfspr Rx, 129 \\
\hline TEXASR & mttexasr Rx & mtspr 130,Rx & mftexasr Rx & mfspr Rx, 130 \\
\hline TEXASRU & mttxasru Rx & mtspr 131,Rx & mftexaru Rx & mfspr Rx, 131 \\
\hline CTRL & - & - & mfctrl Rx & mfspr Rx, 136 \\
\hline VRSAVE & mtvrsave Rx & mtspr 256,Rx & mfvrsave Rx & mfspr Rx,256 \\
\hline SPRG3 & - & - & mfusprg3 Rx & mfspr Rx,259 \\
\hline TB & - & - & mftb Rx & \[
\begin{gathered}
\text { mftb Rx,268 } \\
\text { mfspr Rx,268 }
\end{gathered}
\] \\
\hline TBU & - & - & mftbu Rx & \[
\begin{gathered}
\text { mftb Rx,269 } \\
\text { mfspr Rx,269 }
\end{gathered}
\] \\
\hline SIER & - & - & mfusier Rx & mfspr Rx,768 \\
\hline MMCR2 & mtummcr2 Rx & mtspr 769,Rx & mfummer2 Rx & mfspr Rx,769 \\
\hline MMCRA & mtummera Rx & mtspr 770,Rx & mfummcra Rx & mfspr Rx,770 \\
\hline PMC1 & mtupmc1 Rx & mtspr 771,Rx & mfupmc1 Rx & mfspr Rx,771 \\
\hline PMC2 & mtupmc2 Rx & mtspr 772,Rx & mfupmc2 Rx & mfspr Rx,772 \\
\hline PMC3 & mtupmc3 Rx & mtspr 773,Rx & mfupmc3 Rx & mfspr Rx,773 \\
\hline PMC4 & mtupmc4 Rx & mtspr 774,Rx & mfupmc4 Rx & mfspr Rx,774 \\
\hline PMC5 & mtupmc5 Rx & mtspr 775,Rx & mfupmc5 Rx & mfspr Rx,775 \\
\hline PMC6 & mtupmc6 Rx & mtspr 776,Rx & mfupmc6 Rx & mfspr Rx,776 \\
\hline MMCR0 & mtummcr0 Rx & mtspr 779,Rx & mfummcr0 Rx & mfspr Rx,779 \\
\hline SIAR & - & - & mfusiar Rx & mfspr Rx,780 \\
\hline SDAR & - & - & mfusdar Rx & mfspr Rx,781 \\
\hline MMCR1 & - & - & mfummer1 Rx & mfspr Rx,782 \\
\hline BESCRS & mtbescrs Rx & mtspr 800,Rx & mfbescrs Rx & mfspr Rx, 800 \\
\hline BESCRU & mtbescru Rx & mtspr 801,Rx & mfbescru Rx & mfspr Rx,801 \\
\hline BESCRR & mtbescrr Rx & mtspr 802,Rx & mfbescrr Rx & mfspr Rx,802 \\
\hline BESCRRU & mtbescrru Rx & mtspr 803,Rx & mfbescrru Rx & mfspr Rx,803 \\
\hline EBBHR & mtebbhr Rx & mtspr 804,Rx & mfebbhr Rx & mfspr Rx,804 \\
\hline EBBRR & mtebbrr Rx & mtspr 805,Rx & mfebbrr Rx & mfspr Rx,805 \\
\hline BESCR & mtbescr Rx & mtspr 806,Rx & mfbescr Rx & mfspr Rx,806 \\
\hline TAR & mttar Rx & mtspr 815,Rx & mftar Rx & mfspr Rx,815 \\
\hline PPR & mtppr Rx & mtspr 896,Rx & mfppr Rx & mfspr Rx,896 \\
\hline PPR32 & mtppr32 Rx & mtspr 898,Rx & mfppr32 Rx & mfspr Rx,898 \\
\hline
\end{tabular}

\section*{Version 3.0}

\section*{Examples}
1. Copy the contents of register Rx to the XER.
\(m t x e r \quad\) (equivalent to: mtspr 1,Rx)
2. Copy the contents of the LR to register Rx.
mflr Rx (equivalent to: mfspr Rx,8)
3. Copy the contents of register \(R x\) to the CTR.
mtctr \(R x \quad\) (equivalent to: mtspr 9,Rx)

\section*{C. 10 Miscellaneous Mnemonics}

\section*{No-op}

Many Power ISA instructions can be coded in a way such that, effectively, no operation is performed. An extended mnemonic is provided for the preferred form of no-op. If an implementation performs any type of run-time optimization related to no-ops, the preferred form is the no-op that will trigger this.
\[
\text { nop } \quad \text { (equivalent to: ori } 0,0,0 \text { ) }
\]

For some uses of a no-op instruction, optimizations related to no-ops, such as removal from the execution stream, are not desireable. An extended mnemonic is provided for the executed form of no-op. This form of no-op will still consume execution resources.
xnop (equivalent to: xori \(0,0,0\) )

\section*{Load Immediate}

The addi and addis instructions can be used to load an immediate value into a register. Extended mnemonics are provided to convey the idea that no addition is being performed but merely data movement (from the immediate field of the instruction to a register).
Load a 16-bit signed immediate value into register Rx.
li Rx,value (equivalent to: addi \(R x, 0\), value)
Load a 16-bit signed immediate value, shifted left by 16 bits, into register Rx.
lis \(\quad R x\),value (equivalent to: addis \(R x, 0\), value)

\section*{Load Next Instruction Address}

The addpcis instruction can be used to load the next instruction address into a register. An extended mnemonics is provided to perform this operation.

Inia \(R x \quad\) (equivalent to: addpcis \(R x, 0\) )

\section*{Load Address}

This mnemonic permits computing the value of a base-displacement operand, using the addi instruction which normally requires separate register and immediate operands.
la \(\quad R x, D(R y) \quad\) (equivalent to: addi \(R x, R y, D)\)
The la mnemonic is useful for obtaining the address of a variable specified by name, allowing the Assembler to supply the base register number and compute the displacement. If the variable \(v\) is located at offset Dv bytes from the address in register Rv, and the Assembler has been told to use register Rv as a base for references to the data structure containing \(v\), then the following line causes the address of \(v\) to be loaded into register \(R x\).
la \(\quad \mathrm{Rx}, \mathrm{v} \quad\) (equivalent to: addi \(\mathrm{Rx}, \mathrm{Rv}, \mathrm{Dv}\) )

\section*{Move Register}

Several Power ISA instructions can be coded in a way such that they simply copy the contents of one register to another. An extended mnemonic is provided to convey the idea that no computation is being performed but merely data movement (from one register to another).

The following instruction copies the contents of register Ry to register Rx. This mnemonic can be coded with a final "." to cause the Rc bit to be set in the underlying instruction.
\[
m r \quad R x, R y \quad \text { (equivalent to: or } \quad R x, R y, R y)
\]

\section*{Complement Register}

Several Power ISA instructions can be coded in a way such that they complement the contents of one register and place the result into another register. An extended mnemonic is provided that allows this operation to be coded easily.

The following instruction complements the contents of register Ry and places the result into register Rx. This mnemonic can be coded with a final "." to cause the Rc bit to be set in the underlying instruction.
\[
\text { not } R x, R y \quad \text { (equivalent to: nor } \quad R x, R y, R y \text { ) }
\]

\section*{Move To/From Condition Register}

This mnemonic permits copying the contents of the low-order 32 bits of a GPR to the Condition Register, using the same style as the mfcr instruction.
mtcr Rx (equivalent to: mtcrf \(0 x F F, R x\) )
The following instructions may generate either the (old) mtcrf or mfcr instructions or the (new) mtocrf or mfocrf instruction, respectively, depending on the target machine type assembler parameter.
\begin{tabular}{ll} 
mtcrf & FXM,Rx \\
mfcr & \(R x\)
\end{tabular}

All three extended mnemonics in this subsection are being phased out. In future assemblers the form "mtcr Rx" may not exist, and the mtcrf and mfcr mnemonics may generate the old form instructions (with bit \(11=0\) ) regardless of the target machine type assembler parameter, or may cease to exist.

Version 3.0

\section*{Book II:}

\section*{Power ISA Virtual Environment Architecture}

Version 3.0

\section*{Chapter 1. Storage Model}

\subsection*{1.1 Definitions}

The following definitions, in addition to those specified in Book I, are used in this Book. In these definitions, "Load instruction" includes the Cache Management and other instructions that are stated in the instruction descriptions to be "treated as a Load", and similarly for "Store instruction".

\section*{- system}

A combination of processors, storage, and associated mechanisms that is capable of executing programs. Sometimes the reference to system includes services provided by the privileged software.
- main storage

The level of storage hierarchy in which all storage state is visible to all processors and mechanisms in the system.
- primary cache

The level of cache closest to the processor.
- secondary cache

After the primary cache, the next closest level of cache to the processor.
- instruction storage

The view of storage as seen by the mechanism that fetches instructions.
- data storage

The view of storage as seen by a Load or Store instruction.
- program order

The execution of instructions in the order required by the sequential execution model. (See Section 2.2 of Book I.) A dcbz instruction that modifies storage which contains instructions has the same effect with respect to the sequential execution model as a Store instruction as described there.)
For the instructions and facilities defined in this Book, there are two additional exceptions to the sequential execution model that the processor
obeys beyond those described in Section 2.2 of Book I.
- transaction failure (see Section 5.3.3)
- An event-based branch (see Chapter 7)

\section*{- event-based exception}

An unusual condition, or external signal, that sets a status bit in the BESCR and may or may not cause an event-based branch, depending upon whether event-based branches are enabled.

\section*{- storage location}

A contiguous sequence of one or more bytes in storage. When used in association with a specific instruction or the instruction fetching mechanism, the length of the sequence of one or more bytes is typically implied by the operation. In other uses, it may refer more abstractly to a group of bytes which share common storage attributes.
- storage access

An access to a storage location. There are three (mutually exclusive) kinds of storage access.
- data access

An access to the storage location specified by a Load or Store instruction, or, if the access is performed "out-of-order" (see Section 5.5 of Book III), an access to a storage location as if it were the storage location specified by a Load or Store instruction.
- instruction fetch

An access for the purpose of fetching an instruction.
- implicit access

An access by the processor for the purpose of finding the address translation tables, translating an address, or recording reference and change information (see Book III).

\section*{- caused by, associated with}
- caused by

A storage access is said to be caused by an instruction if the instruction is a Load or Store
and the access (data access) is to the storage location specified by the instruction.

\section*{- associated with}

A storage access is said to be associated with an instruction if the access is for the purpose of fetching the instruction (instruction fetch), or is a data access caused by the instruction, or is an implicit access that occurs as a side effect of fetching or executing the instruction.

■ prefetched instructions
Instructions for which a copy of the instruction has been fetched from instruction storage, but the instruction has not yet been executed.
- uniprocessor

A system that contains one processor.
- multiprocessor

A system that contains two or more processors.
■ shared storage multiprocessor
A multiprocessor that contains some common storage, which all the processors in the system can access.

\section*{- performed}

A load or instruction fetch by a processor or mechanism (P1) is performed with respect to any processor or mechanism (P2) when the value to be returned by the load or instruction fetch can no longer be changed by a store by P2. A store by P1 is performed with respect to P2 when a load by P2 from the location accessed by the store will return the value stored (or a value stored subsequently). An instruction cache block invalidation by P1 is performed with respect to P2 when the instruction that requested the invalidation has caused the specified block, if present, to be made invalid in P2's instruction cache, and similarly for a data cache block invalidation.

The preceding definitions apply regardless of whether P1 and P2 are the same entity.
- page (virtual page)
\(2^{n}\) contiguous bytes of storage aligned such that the effective address of the first byte in the page is an integral multiple of the page size for which protection and control attributes are independently specifiable and for which reference and change status are independently recorded.
- block

The aligned unit of storage operated on by the Cache Management instructions. The size of an instruction cache block may differ from the size of a data cache block, and both sizes may vary between implementations. The maximum block size is equal to the minimum page size.

\section*{- aggregate store}

The set of stores caused by a successful transaction, which are performed as an atomic unit.

\subsection*{1.2 Introduction}

The Power ISA User Instruction Set Architecture, discussed in Book I, defines storage as a linear array of bytes indexed from 0 to a maximum of \(2^{64}-1\). Each byte is identified by its index, called its address, and each byte contains a value. This information is sufficient to allow the programming of applications that require no special features of any particular system environment. The Power ISA Virtual Environment Architecture, described herein, expands this simple storage model to include caches, virtual storage, and shared storage multiprocessors. The Power ISA Virtual Environment Architecture, in conjunction with services based on the Power ISA Operating Environment Architecture (see Book III) and provided by the operating system, permits explicit control of this expanded storage model. A simple model for sequential execution allows at most one storage access to be performed at a time and requires that all storage accesses appear to be performed in program order. In contrast to this simple model, the Power ISA specifies a relaxed model of storage consistency. In a multiprocessor system that allows multiple copies of a storage location, aggressive implementations of the architecture can permit intervals of time during which different copies of a storage location have different values. This chapter describes features of the Power ISA that enable programmers to write correct programs for this storage model.

\subsection*{1.3 Virtual Storage}

The Power ISA system implements a virtual storage model for applications. This means that a combination of hardware and software can present a storage model that allows applications to exist within a "virtual" address space larger than either the effective address space or the real address space.
Each program can access \(2^{64}\) bytes of "effective address" (EA) space, subject to limitations imposed by the operating system. In a typical Power ISA system, each program's EA space is a subset of a larger "virtual address" (VA) space managed by the operating system.

Each effective address is translated to a real address (i.e., to an address of a byte in real storage or on an I/O device) before being used to access storage. The hardware accomplishes this, using the address translation mechanism described in Book III. The operating system manages the real (physical) storage resources of the system, by setting up the tables and other information used by the hardware address translation mechanism.

In general, real storage may not be large enough to map all the virtual pages used by the currently active applications. With support provided by hardware, the operating system can attempt to use the available real pages to map a sufficient set of virtual pages of the applications. If a sufficient set is maintained, "paging" activity is minimized. If not, performance degradation is likely.

The operating system can support restricted access to virtual pages (including read/write, read only, and no access; see Book III), based on system standards (e.g., program code might be read only) and application requests.

\subsection*{1.4 Single-Copy Atomicity}

An access is single-copy atomic, or simply atomic, if it is always performed in its entirety with no visible fragmentation. Atomic accesses are thus serialized: each happens in its entirety in some order, even when that order is not specified in the program or enforced between processors.
The access caused by an instruction other than a Load/ Store Multiple or Move Assist instruction is guaranteed to be atomic if the storage operand is not larger than a doubleword and is aligned (see Section 1.11.1 of Book \(\mathrm{I})\).

Quadword accesses with aligned storage operands are guaranteed to be atomic when caused by the following instructions.
- Iq
- stq
- lqarx
- stqcx.

Quadword atomicity applies only to storage that is neither Write Through Required nor Caching Inhibited. The cases described above are the only cases in which the access to the storage operand is guaranteed to be atomic. For example, the access caused by the following instructions is not guaranteed to be atomic.
- any Load or Store instruction for which the storage operand is unaligned
■ Imw, stmw, Iswi, Iswx, stswi, stswx
- Ifdp, Ifdpx, stfdp, stfdpx
- any Cache Management instruction

An access that is not atomic is performed as a set of smaller disjoint atomic accesses. If the non-atomic access is caused by an instruction other than a Load/ Store Multiple or Move Assist instruction and one of the following conditions is satisfied, the non-atomic access is performed as described in the corresponding list item. The first list item matching a given situation applies.
■ The storage operand is one quadword and is dou-bleword-aligned:
the access is performed as two disjoint aligned doubleword atomic accesses.
- The storage operand is at least eight bytes long and is word-aligned:
the access is performed as a set of disjoint atomic accesses each of which consists of one or more aligned words.
- The storage operand is at least four bytes long and is halfword-aligned:
the access is performed as a set of disjoint atomic accesses each of which consists of one or more aligned halfwords.

In all other cases the number, length, and alignment of the component disjoint atomic accesses are implemen-tation-dependent. In all cases the relative order in which the component disjoint atomic accesses are performed is implementation-dependent.
The results for several combinations of loads and stores to the same or overlapping locations are described below.
1. When two processors perform atomic stores to locations that do not overlap, and no other stores are performed to those locations, the contents of those locations are the same as if the two stores were performed by a single processor.
2. When two processors perform atomic stores to the same storage location, and no other store is performed to that location, the contents of that location are the result stored by one of the processors.
3. When two processors perform stores that have the same target location and are not guaranteed to be atomic, and no other store is performed to that location, the result is some combination of the bytes stored by both processors.
4. When two processors perform stores to overlapping locations, and no other store is performed to those locations, the result is some combination of the bytes stored by the processors to the overlapping bytes. The portions of the locations that do not overlap contain the bytes stored by the processor storing to the location.
5. When a processor performs an atomic store to a location, a second processor performs an atomic load from that location, and no other store is performed to that location, the value returned by the load is the contents of the location before the store or the contents of the location after the store.
6. When a load and a store with the same target location can be performed simultaneously, and the accesses are not guaranteed to be atomic, and no other store is performed to that location, the value returned by the load is some combination of the contents of the location before the store and the contents of the location after the store.

\subsection*{1.5 Cache Model}

A cache model in which there is one cache for instructions and another cache for data is called a "Har-vard-style" cache. This is the model assumed by the Power ISA, e.g., in the descriptions of the Cache Management instructions in Section 4.3. Alternative cache models may be implemented (e.g., a "combined cache" model, in which a single cache is used for both instructions and data, or a model in which there are several levels of caches), but they support the programming model implied by a Harvard-style cache.
The processor is not required to maintain copies of storage locations in the instruction cache consistent with modifications to those storage locations (e.g., modifications caused by Store instructions).
A location in the data cache is considered to be modified in that cache if the location has been modified (e.g., by a Store instruction) and the modified data have not been written to main storage.

Cache Management instructions are provided so that programs can manage the caches when needed. For example, program management of the caches is needed when a program generates or modifies code that will be executed (i.e., when the program modifies data in storage and then attempts to execute the modified data as instructions). The Cache Management instructions are also useful in optimizing the use of memory bandwidth in such applications as graphics and numerically intensive computing. The functions performed by these instructions depend on the storage control attributes associated with the specified storage location (see Section 1.6, "Storage Control Attributes").
The Cache Management instructions allow the program to do the following.
- invalidate the copy of storage in an instruction cache block (icbi)
- provide a hint that an instruction will probably soon be accessed from a specified instruction cache block (icbt)
- provide a hint that the program will probably soon access a specified data cache block (dcbt, dcbtst)
- set the contents of a data cache block to zeros (dcbz)
- copy the contents of a modified data cache block to main storage (dcbst)
- copy the contents of a modified data cache block to main storage and make the copy of the block in the data cache invalid (dcbf or dcbfl)

\subsection*{1.6 Storage Control Attributes}

Some operating systems may provide a means to allow programs to specify the storage control attributes described in this section. Because the support provided for these attributes by the operating system may
vary between systems, the details of the specific system being used must be known before these attributes can be used.

Storage control attributes are associated with units of storage that are multiples of the page size. Each storage access is performed according to the storage control attributes of the specified storage location, as described below. The storage control attributes are the following.
- Write Through Required
- Caching Inhibited
- Memory Coherence Required
- Guarded
- Strong Access Order

These attributes have meaning only when an effective address is translated by the processor performing the storage access.

\section*{Programming Note}

The Write Through Required and Caching Inhibited attributes are mutually exclusive because, as described below, the Write Through Required attribute permits the storage location to be in the data cache while the Caching Inhibited attribute does not.
Storage that is Write Through Required or Caching Inhibited is not intended to be used for general-purpose programming. For example, the Ibarx, Iharx, Iwarx, Idarx, Iqarx, stbcx., sthcx., stwcx., stdcx., and stqcx. instructions may cause the system data storage error handler to be invoked if they specify a location in storage having either of these attributes. To obtain the best performance across the widest range of implementations, storage that is Write Through Required or Caching Inhibited should be used only when the use of such storage meets specific functional or semantic needs or enables a performance optimization.

In the remainder of this section, "Load instruction" includes the Cache Management and other instructions that are stated in the instruction descriptions to be "treated as a Load" unless they are explicitly excluded, and similarly for "Store instruction".

\subsection*{1.6.1 Write Through Required}

A store to a Write Through Required storage location is performed in main storage. A Store instruction that specifies a location in Write Through Required storage may cause additional locations in main storage to be accessed. If a copy of the block containing the specified location is retained in the data cache, the store is also performed in the data cache. The store does not cause the block to be considered to be modified in the data cache.

In general, accesses caused by separate Store instructions that specify locations in Write Through Required storage may be combined into one access. Such combining does not occur if the Store instructions are separated by a sync, eieio instruction.

\subsection*{1.6.2 Caching Inhibited}

An access to a Caching Inhibited storage location is performed in main storage. A Load instruction that specifies a location in Caching Inhibited storage may cause additional locations in main storage to be accessed unless the specified location is also Guarded. An instruction fetch from Caching Inhibited storage may cause additional words in main storage to be accessed. No copy of the accessed locations is placed into the caches.

In general, non-overlapping accesses caused by separate Load instructions that specify locations in Caching Inhibited storage may be combined into one access, as may non-overlapping accesses caused by separate Store instructions that specify locations in Caching Inhibited storage. Such combining does not occur if the Load or Store instructions are separated by a sync instruction. Combining may also occur among such accesses from multiple processors that share a common memory interface. No combining occurs if the storage is also Guarded.

\section*{- Programming Note}

None of the memory barrier instructions prevent the combining of accesses from different processors. The Guarded storage attribute must be used in combination with Caching Inhibited to prevent such combining.

\subsection*{1.6.3 Memory Coherence Required}

An access to a Memory Coherence Required storage location is performed coherently, as follows.
Memory coherence refers to the ordering of stores to a single location. Atomic stores to a given location are coherent if they are serialized in some order, and no processor or mechanism is able to observe any subset of those stores as occurring in a conflicting order. This serialization order is an abstract sequence of values; the physical storage location need not assume each of the values written to it. For example, a processor may update a location several times before the value is written to physical storage. The result of a store operation is not available to every processor or mechanism at the same instant, and it may be that a processor or mechanism observes only some of the values that are written to a location. However, when a location is accessed atomically and coherently by all processors and mechanisms, the sequence of values loaded from the loca-
tion by any processor or mechanism during any interval of time forms a subsequence of the sequence of values that the location logically held during that interval. That is, a processor or mechanism can never load a "newer" value first and then, later, load an "older" value.

Memory coherence is managed in blocks called coherence blocks. Their size is implementation-dependent, but is larger than a word and is usually the size of a cache block.

For storage that is not Memory Coherence Required, software must explicitly manage memory coherence to the extent required by program correctness. The operations required to do this may be system-dependent.
Because the Memory Coherence Required attribute for a given storage location is of little use unless all processors that access the location do so coherently, in statements about Memory Coherence Required storage elsewhere in this document it is generally assumed that the storage has the Memory Coherence Required attribute for all processors that access it.

\section*{Programming Note}

Operating systems that allow programs to request that storage not be Memory Coherence Required should provide services to assist in managing memory coherence for such storage, including all system-dependent aspects thereof.
In most systems the default is that all storage is Memory Coherence Required. For some applications in some systems, software management of coherence may yield better performance. In such cases, a program can request that a given unit of storage not be Memory Coherence Required, and can manage the coherence of that storage by using the sync instruction, the Cache Management instructions, and services provided by the operating system.

\subsection*{1.6.4 Guarded}

A data access to a Guarded storage location is performed only if either (a) the access is caused by an instruction that is known to be required by the sequential execution model, or (b) the access is a load and the storage location is already in a cache. If the storage is also Caching Inhibited, only the storage location specified by the instruction is accessed; otherwise any storage location in the cache block containing the specified storage location may be accessed.

Instructions are not fetched from virtual storage that is Guarded. If the instruction addressed by the current instruction address is in such storage, the system instruction storage error handler may be invoked (see Section 6.5.5 of Book III).

\section*{Version 3.0}

\section*{Programming Note}

In some implementations, instructions may be executed before they are known to be required by the sequential execution model. Because the results of instructions executed in this manner are discarded if it is later determined that those instructions would not have been executed in the sequential execution model, this behavior does not affect most programs.

This behavior does affect programs that access storage locations that are not "well-behaved" (e.g., a storage location that represents a control register on an I/O device that, when accessed, causes the device to perform an operation). To avoid unintended results, programs that access such storage locations should request that the storage be Guarded, and should prevent such storage locations from being in a cache (e.g., by requesting that the storage also be Caching Inhibited).

\subsection*{1.6.5 Strong Access Order}

All accesses to storage with the Strong Access Order (SAO) attribute (referred to as SAO storage) will be performed using a set of ordering rules different from that of the weakly consistent model that is described in Section 1.7.1, "Storage Access Ordering". These rules apply only to accesses that are caused by a Load or a Store, and not to accesses associated with those instructions. Furthermore, these rules do not apply to accesses that are caused by or associated with instructions that are stated in their descriptions to be "treated as a Load" or "treated as a Store." The details are described below, from the programmer's point of view. (The processor may deviate from these rules if the programmer cannot detect the deviation.) The SAO attribute is not intended to be used for general purpose programming. It is provided in a manner that is not fully independent of the other storage attributes. Specifically, it is only provided for storage that is Memory Coherence Required, but not Write Through Required, not Caching Inhibited, and not Guarded. See Section 5.8.2.1, "Storage Control Bit Restrictions", in Book III for more details. Accesses to SAO storage are likely to be performed more slowly than similar accesses to non-SAO storage.

The order in which a processor performs storage accesses to SAO storage, the order in which those accesses are performed with respect to other processors and mechanisms, and the order in which those accesses are performed in main storage are the same except in the circumstances described in the following paragraph. The ordering rules for accesses performed by a single processor to SAO storage are as follows. Stores are performed in program order. When a store accesses data adjacent to that which is accessed by the next store in program order, the two storage accesses may be combined into a single larger access. Loads are performed in program order. When a load accesses data adjacent to that which is accessed by the next load in program order, the two storage accesses may be combined into a single larger access. Stores may not be performed before loads which precede them in program order. Loads may be performed before stores which precede them in program order, with the provision that a load which follows a store of the same datum (to the same address) must obtain a value which is no older (in consideration of the possibility of programs on other processors sharing the same storage) than the value stored by the preceding store.

When any given processor loads the datum it just stored, as described above, the load may be performed by the processor before the preceding store has been performed with respect to other processors and mechanisms, and in main storage. This may cause the processor to see its store earlier relative to stores performed by other processors than it is observed by other processors and mechanisms, and than it is performed in memory. A direct consequence of this con-
sideration is that although programs running on each processor will see the same sequence of accesses from any individual processor to SAO storage, each may in general see a different interleaving of the individual sequences. The memory barrier instructions may be used to establish stronger ordering, as described in Section 1.7.1, "Storage Access Ordering", beginning with the third major bullet.

\subsection*{1.7 Shared Storage}

This architecture supports the sharing of storage between programs, between different instances of the same program, and between processors and other mechanisms. It also supports access to a storage location by one or more programs using different effective addresses. All these cases are considered storage sharing. Storage is shared in blocks that are an integral number of pages.

When the same storage location has different effective addresses, the addresses are said to be aliases. Each application can be granted separate access privileges to aliased pages.

\subsection*{1.7.1 Storage Access Ordering}

The Power ISA defines two models for the ordering of storage accesses: weakly consistent and strong access ordering. The predominant model is weakly consistent. This model provides an opportunity for improved performance over a model that has stronger consistency rules, but places the responsibility on the program to ensure that ordering or synchronization instructions are properly placed when storage is shared by two or more programs. Implementations which support SAO apply a stronger consistency model among accesses to SAO storage. The order between accesses to SAO storage and those performed using the weakly consistent model is characteristic of the weakly consistent model. The following description, through the second major bullet, applies only to the weakly consistent model. The corresponding description for SAO storage is found in Section 1.6.5, "Strong Access Order". The rest of the description following the second bulletted item applies to both models.
The order in which the processor performs storage accesses, the order in which those accesses are performed with respect to another processor or mechanism, and the order in which those accesses are performed in main storage may all be different. Several means of enforcing an ordering of storage accesses are provided to allow programs to share storage with other programs, or with mechanisms such as I/O devices. These means are listed below. The phrase "to the extent required by the associated Memory Coherence Required attributes" refers to the Memory Coherence Required attribute, if any, associated with each access.
- If two Store instructions or two Load instructions specify storage locations that are both Caching Inhibited and Guarded, the corresponding storage accesses are performed in program order with respect to any processor or mechanism.
- If a Load instruction depends on the value returned by a preceding Load instruction (because the value is used to compute the effective address
specified by the second Load), the corresponding storage accesses are performed in program order with respect to any processor or mechanism to the extent required by the associated Memory Coherence Required attributes. This applies even if the dependency has no effect on program logic (e.g., the value returned by the first Load is ANDed with zero and then added to the effective address specified by the second Load).
- When a processor (P1) executes a Synchronize or eieio instruction a memory barrier is created, which orders applicable storage accesses pairwise, as follows. Let \(A\) be a set of storage accesses that includes all storage accesses associated with instructions preceding the barrier-creating instruction, and let \(B\) be a set of storage accesses that includes all storage accesses associated with instructions following the barrier-creating instruction. For each applicable pair \(\mathrm{a}_{\mathrm{i}}, \mathrm{b}_{\mathrm{j}}\) of storage accesses such that \(a_{i}\) is in \(A\) and \(b_{j}\) is in \(B\), the memory barrier ensures that \(a_{i}\) will be performed with respect to any processor or mechanism, to the extent required by the associated Memory Coherence Required attributes, before \(b_{j}\) is performed with respect to that processor or mechanism.

The ordering done by a memory barrier is said to be "cumulative" if it also orders storage accesses that are performed by processors and mechanisms other than P1, as follows.
- A includes all applicable storage accesses by any such processor or mechanism that have been performed with respect to P1 before the memory barrier is created.
- B includes all applicable storage accesses by any such processor or mechanism that are performed after a Load instruction executed by that processor or mechanism has returned the value stored by a store that is in \(B\).

No ordering should be assumed among the storage accesses caused by a single instruction (i.e, by an instruction for which the access is not atomic), even if the accesses are to SAO storage, and no means are provided for controlling that order.

\section*{Programming Note}

Because stores cannot be performed "out-of-order" (see Book III), if a Store instruction depends on the value returned by a preceding Load instruction (because the value returned by the Load is used to compute either the effective address specified by the Store or the value to be stored), the corresponding storage accesses are performed in program order. The same applies if whether the Store instruction is executed depends on a conditional Branch instruction that in turn depends on the value returned by a preceding Load instruction.

Because an isync instruction prevents the execution of instructions following the isync until instructions preceding the isync have completed, if an isync follows a conditional Branch instruction that depends on the value returned by a preceding Load instruction, the load on which the Branch depends is performed before any loads caused by instructions following the isync. This applies even if the effects of the "dependency" are independent of the value loaded (e.g., the value is compared to itself and the Branch tests the EQ bit in the selected CR field), and even if the branch target is the sequentially next instruction.
With the exception of the cases described above and earlier in this section, data dependencies and control dependencies do not order storage accesses. Examples include the following.
- If a Load instruction specifies the same storage location as a preceding Store instruction and the location is in storage that is not Caching Inhibited, the load may be satisfied from a "store queue" (a buffer into which the processor places stored values before presenting them to the storage subsystem), and not be visible to other processors and mechanisms. A consequence is that if a subsequent Store depends on the value returned by the Load, the two stores need not be performed in program order with respect to other processors and mechanisms.
- Because a Store Conditional instruction may complete before its store has been performed, a conditional Branch instruction that depends on the CRO value set by a Store Conditional instruction does
not order the Store Conditional's store with respect to storage accesses caused by instructions that follow the Branch.

■ Because processors may predict branch target addresses and branch condition resolution, control dependencies (e.g., branches) do not order storage accesses except as described above. For example, when a subroutine returns to its caller the return address may be predicted, with the result that loads caused by instructions at or after the return address may be performed before the load that obtains the return address is performed.
Because processors may implement nonarchitected duplicates of architected resources (e.g., GPRs, CR fields, and the Link Register), resource dependencies (e.g., specification of the same target register for two Load instructions) do not order storage accesses.

Examples of correct uses of dependencies, sync and Iwsync to order storage accesses can be found in Appendix B. "Programming Examples for Sharing Storage" on page 913.
Because the storage model is weakly consistent, the sequential execution model as applied to instructions that cause storage accesses guarantees only that those accesses appear to be performed in program order with respect to the processor executing the instructions. For example, an instruction may complete, and subsequent instructions may be executed, before storage accesses caused by the first instruction have been performed. However, for a sequence of atomic accesses to the same storage location, if the location is in storage that is Memory Coherence Required the definition of coherence guarantees that the accesses are performed in program order with respect to any processor or mechanism that accesses the location coherently, and similarly if the location is in storage that is Caching Inhibited.
Because accesses to storage that is Caching Inhibited are performed in main storage, memory barriers and dependencies on Load instructions order such accesses with respect to any processor or mechanism even if the storage is not Memory Coherence Required.

\section*{Programming Note}

The first example below illustrates cumulative ordering of storage accesses preceding a memory barrier, and the second illustrates cumulative ordering of storage accesses following a memory barrier. Assume that locations X, Y, and Z initially contain the value 0 .

\section*{Example 1:}

Processor A:
stores the value 1 to location \(X\)
Processor B:
loads from location X obtaining the value 1, executes a sync instruction, then stores the value 2 to location \(Y\)

Processor C:
loads from location Y obtaining the value 2 , executes a sync instruction, then loads from location \(X\)

\section*{Example 2:}

\section*{Processor A:}
stores the value 1 to location X , executes a sync instruction, then stores the value 2 to location Y

Processor B:
loops loading from location \(Y\) until the value 2 is obtained, then stores the value 3 to location Z

Processor C:
loads from location \(Z\) obtaining the value 3 , executes a sync instruction, then loads from location \(X\)

In both cases, cumulative ordering dictates that the value loaded from location \(X\) by processor \(C\) is 1 .

\subsection*{1.7.2 Storage Ordering of Copy/ Paste-Initiated Data Transfers}

The Copy/Paste Facility (see Section 4.4) uses pairs of instructions to initiate 128-byte data transfers. They are referred to as "data transfers" to differentiate them from the "normal" storage accesses caused by or associated with loads, stores, and instructions that are treated as loads and stores. The facility provides for well-defined "move groups" of such transfers. The relative order among the individual data transfers within a move group cannot be controlled by any means. Moreover, the implicit ordering characteristics of the sequential execution model and of coherence-required storage do not apply to the individual data transfers within a move group. Similarly, in the absence of barriers, the relative ordering among adjacent move groups or move groups and storage accesses is not defined, and the SEM and coherence-required ordering relationships do not apply.

To establish order between adjacent move groups or between move groups and storage accesses, hwsync must be used. More specifically, the execution synchronization performed by hwsync together with the status gathering of paste_last will have the net effect that a move group preceding an hwsync must complete before instructions that follow the hwsync are initiated. See the description of the Synchronize instruction in Section 4.6.3 for more information.

\section*{- Programming Note}

It may be helpful to think of a copy/paste pair sending the real storage addresses of the 128-byte source and destination to an asynchronous data transfer engine completely separate from the processor that is executing the copy and paste instructions. The transfers in a move group collect in the engine's queue. The engine may perform the data transfers in the group in any order, and with the only relative timing relationship to adjacent groups and accesses being determined by hwsync.

\subsection*{1.7.3 Storage Ordering of I/O Accesses}

A "coherence domain" consists of all processors and all interfaces to main storage. Memory reads and writes initiated by mechanisms outside the coherence domain are performed within the coherence domain in the order in which they enter the coherence domain and are performed as coherent accesses.

\subsection*{1.7.4 Atomic Update}

The Load And Reserve and Store Conditional instructions together permit atomic update of a shared storage location. There are byte, halfword, word, doubleword, and quadword forms of each of these instructions. Described here is the operation of the word forms Iwarx and stwcx.; operation of the byte, halfword, doubleword, and quadword forms Ibarx, stbcx., Iharx, sthcx., Idarx, stdcx., Iqarx, and stqcx. is the same except for obvious substitutions.

The Iwarx instruction is a load from a word-aligned location that has two side effects. Both of these side effects occur at the same time that the load is performed.
1. A reservation for a subsequent stwcx. instruction is created.
2. The memory coherence mechanism is notified that a reservation exists for the storage location specified by the Iwarx.
The stwcx. instruction is a store to a word-aligned location that is conditioned on the existence of the reservation created by the Iwarx and on whether the same storage location is specified by both instructions. To emulate an atomic operation with these instructions, it is necessary that both the Iwarx and the stwcx. specify the same storage location.

A stwcx. performs a store to the target storage location only if the storage location specified by the Iwarx that established the reservation has not been stored into by another processor or mechanism since the reservation was created. If the storage locations specified by the two instructions are in different reservation granules, the store is not performed.

A stwcx. that performs its store is said to "succeed".
Examples of the use of Iwarx and stwcx. are given in Appendix B. "Programming Examples for Sharing Storage" on page 913.

A successful stwcx. to a given location may complete before its store has been performed with respect to other processors and mechanisms. As a result, a subsequent load or Iwarx from the given location by another processor may return a "stale" value. However, a subsequent Iwarx from the given location by the other processor followed by a successful stwcx. by that processor is guaranteed to have returned the value stored by the first processor's stwcx. (in the absence of other stores to the given location).

If a Store Conditional instruction is used with a preceding Load and Reserve instruction that has a different storage operand length (e.g., stwcx. with Idarx), the reservation is cleared and it is undefined whether the store is performed.

\section*{Programming Note}

The store caused by a successful stwcx. is ordered, by a dependence on the reservation, with respect to the load caused by the Iwarx that established the reservation, such that the two storage accesses are performed in program order with respect to any processor or mechanism.

\section*{Programming Note}

There is no mechanism for the Store Conditional instruction to detect that a virtual page has been moved to a new real page and back again to the original real page that was accessed by a Load and Reserve instruction. Privileged software that moves a virtual page could clear the reservation on the processor it is running on in order to ensure that a Store Conditional instruction executed by that processor does not succeed in this case. (The stores that occur naturally as part of moving the virtual page will cause any reservations, held by other processors, in the target real page to be lost.)

\subsection*{1.7.4.1 Reservations}

The ability to emulate an atomic operation using Iwarx and stwcx. is based on the conditional behavior of stwcx., the reservation created by Iwarx, and the clearing of that reservation if the target storage location is modified by another processor or mechanism before the stwcx. performs its store.

A reservation is held on an aligned unit of real storage called a reservation granule. The size of the reservation granule is \(2^{n}\) bytes, where \(n\) is implementation-dependent but is always at least 4 (thus the minimum reservation granule size is a quadword) and, where \(2^{n}\) is not larger than the smallest real page size supported by the implementation. The reservation granule associated with effective address EA contains the real address to which EA maps. ("real_addr(EA)" in the RTL for the Load And Reserve and Store Conditional instructions stands for "real address to which EA maps".) The reservation also has an associated length, which is equal to the storage operand length, in bytes, of the Load and Reserve instruction that established the reservation.

A processor has at most one reservation at any time. A reservation is established by executing a Ibarx, Iharx, Iwarx, Idarx, or Iqarx instruction, as described in item 1 below, and is lost or may be lost, depending on the item, if any of the following occur. Items 1-9 apply only if the relevant access is performed. (For example, an access that would ordinarily be caused by an instruction might not be performed if the instruction causes the system error handler to be invoked.)
1. The processor holding the reservation executes another Ibarx, Iharx, Iwarx, or Idarx: this clears the first reservation and establishes a new one.
2. The processor holding the reservation executes any stbcx., sthcx., stwcx., stdcx., or stqcx., regardless of whether the specified address matches the address specified by the Ibarx, Iharx, Iwarx, Idarx, or Iqarx that established the reservation, and regardless of whether the storage operand lengths of the two instructions are the same.
3. The processor holding the reservation executes an AMO that updates the same reservation granule: whether the reservation is lost is undefined.
4. Any of the following occurs on the processor holding the reservation.
a. The transaction state changes (from Non-transactional, Transactional, or Suspended state to one of the other two states; see Section 5.2, "Transactional Memory Facility States"), except in the following cases
- If the change is from Transactional state to Suspended state, the reservation is not lost.
- If the change is from Suspended state to Transactional state, the reservation is not lost if it was established in Transactional state.
- If the change is caused by a treclaim. or trechkpt. instruction, whether the reservation is lost is undefined.
b. The transaction nesting depth (see Section 5.4, "Transactional Memory Facility Registers") changes; whether the reservation is lost is undefined. (This item applies only if the processor is in Transactional state both before and after the change.)
c. The processor is in Suspended state and executes a Store Conditional instruction (stbcx., sthcx., stwcx., stdcx., or stqcx.) or a waitrsv instruction; the reservation is lost if it was established in Transactional state. In this case the Store Conditional instruction's store is not performed, and the waitrsv does not wait. (For Store Conditional, the reservation is also lost if it was established in Suspended state; see item 2.)
5. Some other processor executes a Store or dcbz that specifies a location in the same reservation granule.
6. Some other processor executes a dcbtst, or dcbt that specifies a location in the same reservation granule: whether the reservation is lost is undefined. (For a dcbtst instruction that specifies a data stream, "location" in the preceding sentence includes all locations in the data stream.)
7. Any processor modifies a Reference or Change bit (see Book III in the same reservation granule: whether the reservation is lost is undefined.
8. Some mechanism other than a processor modifies a storage location in the same reservation granule.
9. An interrupt (see Book III) occurs on the processor holding the reservation: the reservation is not lost. However, system software invoked by interrupts may clear the reservation.)
10. Implementation-specific characteristics of the coherence mechanism cause the reservation to be lost.

\section*{Virtualized Implementation Note}

A reservation may be lost if:
- Software executes a privileged instruction or utilizes a privileged facility
- Software accesses storage not intended for general-purpose programming
- Software accesses a Device Control Register

\section*{Programming Note}

One use of Iwarx and stwcx. is to emulate a "Compare and Swap" primitive like that provided by the IBM System/370 Compare and Swap instruction; see Section B.1, "Atomic Update Primitives" on page 913. A System/370-style Compare and Swap checks only that the old and current values of the word being tested are equal, with the result that programs that use such a Compare and Swap to control a shared resource can err if the word has been modified and the old value subsequently restored. The combination of Iwarx and stwcx. improves on such a Compare and Swap, because the reservation reliably binds the Iwarx and stwcx. together. The reservation is always lost if the word is modified by another processor or mechanism between the Iwarx and stwcx., so the stwcx. never succeeds unless the word has not been stored into (by another processor or mechanism) since the Iwarx.

\section*{Programming Note}

In general, programming conventions must ensure that Iwarx and stwcx. specify addresses that match; a stwcx. should be paired with a specific Iwarx to the same storage location. Situations in which a stwcx. may erroneously be issued after some Iwarx other than that with which it is intended to be paired must be scrupulously avoided. For example, there must not be a context switch in which the processor holds a reservation in behalf of the old context, and the new context resumes after a Iwarx and before the paired stwcx. The stwcx. in the new context might succeed, which is not what was intended by the programmer. Such a situation must be prevented by executing a stbcx., sthcx., stwcx., stdcx., or stqcx. that specifies a dummy writable aligned location as part of the context switch; see Section 6.4.3 of Book III.

\section*{- Programming Note}

Because the reservation is lost if another processor stores anywhere in the reservation granule, lock words (or bytes, halfwords, or doublewords) should be allocated such that few such stores occur, other than perhaps to the lock word itself. (Stores by other processors to the lock word result from contention for the lock, and are an expected consequence of using locks to control access to shared storage; stores to other locations in the reservation granule can cause needless reservation loss.) Such allocation can most easily be accomplished by allocating an entire reservation granule for the lock and wasting all but one word. Because reservation granule size is implementation-dependent, portable code must do such allocation dynamically.

Similar considerations apply to other data that are shared directly using Iwarx and stwcx. (e.g., pointers in certain linked lists; see Section B.3, "List Insertion" on page 917).

\subsection*{1.7.4.2 Forward Progress}

Forward progress in loops that use Iwarx and stwcx. is achieved by a cooperative effort among hardware, system software, and application software.
The architecture guarantees that when a processor executes a Iwarx to obtain a reservation for location \(X\) and then a stwcx. to store a value to location X, either
1. the stwcx. succeeds and the value is written to location X, or
2. the stwcx. fails because some other processor or mechanism modified location X, or
3. the stwcx. fails because the processor's reservation was lost for some other reason.

In Cases 1 and 2, the system as a whole makes progress in the sense that some processor successfully modifies location X. Case 3 covers reservation loss required for correct operation of the rest of the system. This includes cancellation caused by some other processor or mechanism writing elsewhere in the reservation granule, cancellation caused by the operating system in managing certain limited resources such as real storage, and cancellation caused by any of the other effects listed in see Section 1.7.4.1.

An implementation may make a forward progress guarantee, defining the conditions under which the system as a whole makes progress. Such a guarantee must specify the possible causes of reservation loss in Case 3. While the architecture alone cannot provide such a guarantee, the characteristics listed in Cases 1 and 2 are necessary conditions for any forward progress guarantee. An implementation and operating system can build on them to provide such a guarantee.

\section*{Virtualized Implementation Note}

On a virtualized implementation, Case 3 includes reservation loss caused by the virtualization software. Thus, on a virtualized implementation, a reservation may be lost at any time without apparent cause. The virtualization software participates in any forward progress assurances, as described above.

\section*{Programming Note}

The architecture does not include a "fairness guarantee". In competing for a reservation, two processors can indefinitely lock out a third.

\subsection*{1.8 Transactions}

A transaction is a group of instructions that collectively have unique storage access behavior intended to facilitate parallel programming. (It is possible to nest transactions within one another. The description in this chapter will ignore nesting because it does not have a significant impact on the properties of the memory model. Nesting and its consequences will be described elsewhere.) Sequences of instructions that are part of the transaction may be interleaved with sequences of Suspended state instructions that are not part of the transaction. A transaction is said to "succeed" or to "fail," and failure may happen before all of the instructions in the transaction have completed. If the transaction fails, it is as if the instructions that are part of the transaction were never executed. If the transaction succeeds, it appears to execute as an atomic unit as viewed by other processors and mechanisms. (Although the transaction appears to execute atomically, some knowledge of the inner workings will be necessary to avoid apparent paradoxes in the rest of the model. These details are described below.) The execution of Suspended state sequences have the same effect that the sequence would have in the absence of a transaction, independent of the success or failure of the transaction, including accessing storage according to the weakly consistent storage model or SAO, based on storage attributes. Upon failure, normal execution continues at the failure handler. Except for the rollback of the effects of transactional instructions upon transaction failure, as viewed by the executing thread, the interleaved sequences of Transactional and Suspended state instructions appear to execute according to the sequential execution model. See Chapter 5. "Transactional Memory Facility" on page 881 for more details. The unique attributes of the storage model for transactions are described below.
Transaction processing does not support the rollback of operations on the reservation mechanism. To prevent this possibility, a reservation is lost as a result of a state change from Transactional to Non-transactional or

Non-transactional to Transactional. It is possible to successfully complete an atomic update in Transactional state, though such a sequence would have no benefit. It is also possible to complete an atomic update in Suspended state, or straddling an interval in Suspended state if Suspended state is entered via an interrupt or tsuspend. and exited via tresume., rfebb,

\section*{I} update will not succeed if only one of the Load and Reserve / Store Conditional instruction pair is executed in Suspended state.

\section*{Programming Note}

Note that if a Store Conditional instruction within a transaction does not store, it may still be possible for the transaction to succeed. Software must not depend on the two operations having the same outcome. For example, software must not use success of an enclosing transaction as a replacement for checking the condition code from a transactional Store Conditional instruction.

\section*{Programming Note}

Accessing storage locations in Suspended state that have been accessed transactionally has the potential to create apparent storage paradoxes. Consider, for example, a case where variable X has intial value zero, is updated transactionally to one, is read in Suspended state, subsequently the transaction fails, and variable \(X\) is read again. In the absence of external conflicts, the observed sequence of values will be zero, one, zero: old, new, old.
Performing an atomic update on \(X\) in Suspended state may be even more confusing. Suppose the atomic sequence increments X , but that the only way to have \(X=1\) is via the transactional store that occurs before entering Suspended state. The store conditional, if it succeeds, will store \(\mathrm{X}=2\) and in so doing, kill the transaction. But with the transaction having failed, X was never equal to one.
The flexibility of the Suspended state programming model can create unintuitive results. It must be used with care.

Successful transactions are serialized in some order, and no processor or mechanism is able to observe the accesses caused by any subset of these transactions as occurring in an order that conflicts with this order. Specifically, let processor i execute transactions \(0,1, \ldots\), \(\mathrm{j}, \mathrm{j}+1, \ldots\), where only successful transactions are numbered, and the numbering reflects program order. Let \(\mathrm{T}_{\mathrm{ij}}\) be transaction j on processor i . Then there is an ordering of the \(\mathrm{T}_{\mathrm{ij}}\) such that no processor or mechanism is able to observe the accesses caused by the transactions \(\mathrm{T}_{\mathrm{ij}}\) in an order that conflicts with this ordering. Note that Suspended state storage accesses are not included in the serialization property.

\section*{- Programming Note \\ The ordering of the \(\mathrm{T}_{\mathrm{ij}}\) for a given i is consistent with program order for processor i.}

Because of the difference between a transaction's instantaneous appearance and the finite time required to execute it in an implementation, it is exposed to changes in memory management state in a way that is not true for individual accesses. A change to the translation or protection state that would prevent any access from taking place at any time during its processing for the transaction compromises the integrity of the transaction. Any such change must either be prevented or must cause the transaction to fail. The architecture will automatically fail a transaction if the memory management state change is accomplished using tlbie or slbieg. An implementation may overdetect such conflicts between the tlbie or slbieg and the transaction footprint. (Overdetection may result from the technique used to detect the conflict. A bloom filter may be used, as an example. Subsequent references to translation invalidation conflicts implicitly include any cases of spurious overdetection.) Changes made in some other manner must be managed by software, for example by explicitly terminating any affected transactions. Examples of instructions that require software management are tlbiel, slbie, and slbia.
The atomic nature of a transaction, together with the cumulative memory barrier created by the transaction and the memory barriers created by tbegin. and tend. described below, has the potential to eliminate the need for explicit memory barriers within the transaction, and before and after the transaction as well. However, since there may be a desire to preserve existing algorithms while exploiting transactions, the interaction of memory barriers and transactions is defined. In the presence of transactions, storage access ordering is the same as if no transactions are present, with the following exceptions. Memory barriers that are created while the transaction is running (other than the integrated cumulative memory barrier of the transaction described below), data dependencies, and SAO do not order transactional stores. Instead, transactional stores are grouped together into an "aggregate store," which is performed as an atomic unit with respect to other processors and mechanisms when the transaction succeeds, after all the transactional loads have been performed. With this store behavior, the appearance of transactional atomicity is created in a manner similarly to that for a Load and Reserve / Store Conditional pair. Success of the transaction is conditional on the storage locations specified by the loads not having been stored into by a more recent Suspended state store or by any store by another processor or mechanism since the load was performed. (There are additional conditions for the success of transactions.)
A tbegin. instruction that begins a successful transaction creates a memory barrier that immediately pre-
cedes the transaction and orders storage accesses pairwise, as follows. Let \(A\) and \(B\) be sets of storage accesses as defined below. For each pair \(a_{i} b_{j}\) of storage accesses such that \(a_{i}\) is in \(A\) and \(b_{j}\) is in \(B\), the memory barrier ensures that \(a_{i}\) will be performed with respect to any processor or mechanism, to the extent required by the associated Memory Coherence Required attributes, before \(b_{j}\) is performed with respect to that processor or mechanism. Set A contains all data accesses caused by instructions preceding the tbegin. that are neither Write Through Required nor Caching Inhibited. Set B contains all data accesses caused by instructions following the tbegin., including Suspended state accesses, that are neither Write Through Required nor Caching Inhibited. The ordering done by this memory barrier is cumulative.

\section*{Programming Note}

The reason the creation of the memory barrier by tbegin. is specified to be contingent on the transaction succeeding is that delaying the creation may improve performance, and does not seriously inconvenience software.

A successful transaction has an integrated memory barrier behavior. When a processor (P1) executes a tend. instruction and tend. processing determines that the transaction will succeed, a memory barrier is created, which orders storage accesses pairwise, as follows. Let \(A\) and \(B\) be sets of storage accesses as defined below. For each pair \(a_{i} b_{j}\) of storage accesses such that \(a_{i}\) is in \(A\) and \(b_{j}\) is in \(B\), the memory barrier ensures that \(a_{i}\) will be performed with respect to any processor or mechanism, to the extent required by the associated Memory Coherence Required attributes, before \(b_{j}\) is performed with respect to that processor or mechanism. Set A contains all non-transactional data accesses by other processors and mechanisms that have been performed with respect to P1 before the memory barrier is created and are neither Write Through Required nor Caching Inhibited. Set B contains the aggregate store and all non-transactional data accesses by other processors and mechanisms that are performed after a Load instruction executed by that processor or mechanism has returned the value stored by a store that is in set B. Note that the cumulative memory barrier does not order Suspended state storage accesses interleaved with the transaction. The ordering done by this memory barrier is cumulative.
A tend. instruction that ends a successful transaction creates a memory barrier that immediately follows the transaction and orders storage accesses pairwise, as follows. Let A and B be sets of storage accesses as defined below. For each pair \(a_{i} b_{j}\) of storage accesses such that \(a_{i}\) is in \(A\) and \(b_{j}\) is in \(B\), the memory barrier ensures that \(a_{i}\) will be performed with respect to any processor or mechanism, to the extent required by the associated Memory Coherence Required attributes, before \(b_{j}\) is performed with respect to that processor or
mechanism. Set A contains all data accesses caused by instructions preceding the tend., including Suspended state accesses, that are neither Write Through Required nor Caching Inhibited. Set B contains all data accesses caused by instructions following the tend. that are neither Write Through Required nor Caching Inhibited.

\section*{Programming Note}

The memory barriers that are created by the execution of a successful transaction (those associated with tbegin., tend., and the integrated cumulative memomry barrier) render most explicit memory barriers in and around transactions redundant. An exception is when there is a need to establish order among Suspended state accesses.

\subsection*{1.8.1 Rollback-Only Transactions}

A Rollback-Only Transaction (ROT) is a sequence of instructions that is executed, or not, as a unit. The purpose of the ROT is to enable bulk speculation of instructions with minimum overhead. It leverages the rollback mechanism that is invoked as part of transaction failure handling, but has reduced overhead in that it does not have the full atomic nature of the transaction and its synchronization and serialization properties. The absence of a (normal) transaction's atomic quality means that a ROT must not be used to manipulate shared data.

More specifically, a ROT differs from a normal transaction as follows.
- ROTs are not serialized.
- There are no memory barriers created by tbegin. and tend.
- A ROT has no integrated cumulative memory barrier.
■ There is no monitoring of storage locations specified by loads for modification by other processors and mechanisms between the performing of the loads and the completion of the ROT.
■ The stores that are included in the ROT need not appear to be performed as an aggregate store. (Implementations are likely to provide an aggregate store appearance, but the correctness of the program must not depend on the aggregate store appearance.)

\subsection*{1.9 Cluster Shared Memory}

Computer systems may be connected to provide a means for one to directly address storage in another. The group of systems so interconnected is called a "cluster." Real address space in each system in the cluster may be designated to map to "Cluster Shared Memory" or CSM. Main storage from all systems in the
cluster may be mapped into CSM. The mapping of real address space to CSM, the mapping of CSM to main storage, and the means, if any, by which software can control the mappings, are implementation-specific.

The access semantics for CSM are different from the semantics for non-CSM address space. CSM may only be accessed via the Copy-Paste Facility. Furthermore, for any given move involving CSM, the other end of the move must be a line of \(I=0\) storage that is located on the system executing the move. Access via loads, stores, accesses that are treated as loads or stores, and instruction fetching is prohibited. Other semantics are described in Section 1.7.2, "Storage Ordering of Copy/Paste-Initiated Data Transfers".

\subsection*{1.10 Instruction Storage}

The instruction execution properties and requirements described in this section, including its subsections, apply only to instruction execution that is required by the sequential execution model.

In this section, including its subsections, it is assumed that all instructions for which execution is attempted are in storage that is not Caching Inhibited and (unless instruction address translation is disabled; see Book III) is not Guarded, and from which instruction fetching does not cause the system error handler to be invoked (e.g., from which instruction fetching is not prohibited by the "address translation mechanism" or the "storage protection mechanism"; see Book III).

\section*{Programming Note}

The results of attempting to execute instructions from storage that does not satisfy this assumption are described in Section 1.6.2 and Section 1.6.4 of this Book and in Book III.

For each instance of executing an instruction from location X , the instruction may be fetched multiple times.

The instruction cache is not necessarily kept consistent with the data cache or with main storage. It is the responsibility of software to ensure that instruction storage is consistent with data storage when such consistency is required for program correctness.
After one or more bytes of a storage location have been modified and before an instruction located in that storage location is executed, software must execute the appropriate sequence of instructions to make instruction storage consistent with data storage. Otherwise the result of attempting to execute the instruction is boundedly undefined except as described in Section 1.10.1, "Concurrent Modification and Execution of Instructions" on page 828.

\section*{Programming Note}

Following are examples of how to make instruction storage consistent with data storage. Because the optimal instruction sequence to make instruction storage consistent with data storage may vary between systems, many operating systems will provide a system service to perform this function.

Case 1: The given program does not modify instructions executed by another program nor does another program modify the instructions executed by the given program.

Assume that location \(X\) previously contained the instruction AO; the program modified one of more bytes of that location such that, in data storage, the location contains the instruction A1; and location X is wholly contained in a single cache block. The following instruction sequence will make instruction storage consistent with data storage such that if the isync was in location \(\mathrm{X}-4\), the instruction A1 in location X would be executed immediately after the isync.
\begin{tabular}{lll} 
dcbst \(X\) & \#copy the block to main storage \\
sync & \#order copy before invalidation \\
icbi & \(X\) & \#invalidate copy in instr cache \\
isync & & \#discard prefetched instructions
\end{tabular}

Case 2: One or more programs execute the instructions that are concurrently being modified by another program.

Assume program A has modified the instruction at location X and other programs are waiting for program A to signal that the new instruction is ready to execute. The following instruction sequence will make instruction storage consistent with data storage and then set a flag to indicate to the waiting programs that the new instruction can be executed.
\begin{tabular}{lll} 
li r0,1 & \#put a 1 value in r0 \\
dcbst X & \begin{tabular}{l} 
\#copy the block in main storage \\
\#order copy before invalidation
\end{tabular} \\
sync & X & \begin{tabular}{l} 
\#invalidate copy in instr cache \\
sync
\end{tabular} \\
\#order invalidation before store \\
\# to flag
\end{tabular}

The following instruction sequence, executed by the waiting program, will prevent the waiting programs from executing the instruction at location \(X\) until location \(X\) in instruction storage is consistent with data storage, and then will cause any prefetched instructions to be discarded.
\begin{tabular}{llll} 
lwz & r0, flag & \#loop until flag \(=1\) (when 1 is \\
cmpwi & r0,1 & \(\#\) loaded, location \(X\) in inst'n \\
bne & \(\$-8\) & \(\#\) storage is consistent with \\
& & \(\#\) location X in data storage) \\
isync & & \# \#iscard any prefetched inst'ns
\end{tabular}

In the preceding instruction sequence any context synchronizing instruction (e.g., rfid) can be used instead of isync. (For Case 1 only isync can be used.)
For both cases, if two or more instructions in separate data cache blocks have been modified, the dcbst instruction in the examples must be replaced by a sequence of dcbst instructions such that each block containing the modified instructions is copied back to main storage. Similarly, for icbi the sequence must invalidate each instruction cache block containing a location of an instruction that was modified. The sync instruction that appears above between "dcbst \(X\) " and "icbi X" would be placed between the sequence of dcbst instructions and the sequence of icbi instructions.

\subsection*{1.10.1 Concurrent Modification and Execution of Instructions}

The phrase "concurrent modification and execution of instructions" (CMODX) refers to the case in which a processor fetches and executes an instruction from instruction storage which is not consistent with data storage or which becomes inconsistent with data storage prior to the completion of its processing. This section describes the only case in which executing this instruction under these conditions produces defined results.

In the remainder of this section the following terminology is used.
- Location X is an arbitrary word-aligned storage location.
- \(X_{0}\) is the value of the contents of location \(X\) for which software has made the location \(X\) in instruction storage consistent with data storage.
■ \(X_{1}, X_{2}, \ldots, X_{n}\) are the sequence of the first \(n\) values occupying location \(X\) after \(X_{0}\).
- \(X_{n}\) is the first value of \(X\) subsequent to \(X_{0}\) for which software has again made instruction storage consistent with data storage.
■ The "patch class" of instructions consists of the l-form Branch instruction ( \(\boldsymbol{b}[\boldsymbol{I}[\mathbf{a}]\) ) and the preferred no-op instruction (ori \(0,0,0\) ).

If the instruction from location \(X\) is executed after the copy of location \(X\) in instruction storage is made consistent for the value \(X_{0}\) and before it is made consistent for the value \(X_{n}\), the results of executing the instruction are defined if and only if the following conditions are satisfied.
1. The stores that place the values \(X_{1}, \ldots, X_{n}\) into location \(X\) are atomic stores that modify all four bytes of location X .
2. Each \(\mathrm{X}_{\mathrm{i}}, 0 \leq \mathrm{i} \leq \mathrm{n}\), is a patch class instruction.
3. Location \(X\) is in storage that is Memory Coherence Required.

If these conditions are satisfied, the result of each execution of an instruction from location \(X\) will be the execution of some \(X_{i}, 0 \leq i \leq n\). The value of the ordinate \(i\) associated with each value executed may be different and the sequence of ordinates i associated with a sequence of values executed is not constrained, (e.g., a valid sequence of executions of the instruction at location \(X\) could be the sequence \(X_{i}, X_{i+2}\), then \(X_{i-1}\) ). If these conditions are not satisfied, the results of each such execution of an instruction from location \(X\) are boundedly undefined, and may include causing inconsistent information to be presented to the system error handler.

\section*{Programming Note}

An example of how failure to satisfy the requirements given above can cause inconsistent information to be presented to the system error handler is as follows. If the value \(X_{0}\) (an illegal instruction) is executed, causing the system illegal instruction handler to be invoked, and before the error handler can load \(X_{0}\) into a register, \(X_{0}\) is replaced with \(X_{1}\), an Add Immediate instruction, it will appear that a legal instruction caused an illegal instruction exception.

\section*{Programming Note}

It is possible to apply a patch or to instrument a given program without the need to suspend or halt the program. This can be accomplished by modifying the example shown in the Programming Note at the end of Section 1.10 where one program is creating instructions to be executed by one or more other programs.
In place of the Store to a flag to indicate to the other programs that the code is ready to be executed, the program that is applying the patch would replace a patch class instruction in the original program with a Branch instruction that would cause any program executing the Branch to branch to the newly created code. The first instruction in the newly created code must be an isync, which will cause any prefetched instructions to be discarded, ensuring that the execution is consistent with the newly created code. The instruction storage location containing the isync instruction in the patch area must be consistent with data storage with respect to the processor that will execute the patched code before the Store which stores the new Branch instruction is performed.

\section*{Programming Note}

It is believed that all processors that comply with versions of the architecture that precede Version 2.01 support concurrent modification and execution of instructions as described in this section if the requirements given above are satisfied, and that most such processors yield boundedly undefined results if the requirements given above are not satisfied. However, in general such support has not been verified by processor testing. Also, one such processor is known to yield undefined results in certain cases if the requirements given above are not satisfied.

\title{
Chapter 2. Performance Considerations and Instruction Restart
}

\subsection*{2.1 Performance-Optimized Instruction Sequences}

Performance-optimized instruction sequences are instruction sequences that provide better performance than other ways of achieving the same results. The supported performance-optimized sequences are
| shown in the following sections. In order to achieve the improved performance, the sequences must be coded exactly as shown, including instruction order, register
I re-use, and lack of intervening instructions. The processor achieves the improved performance by executing the sequence as a single operation, or in some other highly efficient, sequence-specific, manner. (The improved performance may not be obtained if the sequence causes the system error handler to be invoked, or for implementation-dependent reasons.)
I

\section*{Version 3.0}

\subsection*{2.1.1 Load and Store Operations}

The following instruction sequences will optimize performance for storage accesses to effective addresses that are offset from (RA) by magnitudes of up to \(2^{32}\).
\begin{tabular}{l} 
Operation \\
\hline Fixed-point byte accesses \\
\hline Fixed-point halfword accesses \\
\hline Sequence
\end{tabular}

Table 1: Loads and Stores with offsets of up to \(2^{32}\) offsets from base register

The following instruction sequences will optimize performance for storage accesses to effective addresses that are offset from (RB) by magnitudes of up to \(2^{16}\).
\begin{tabular}{|c|c|c|c|c|}
\hline Operation & \multicolumn{2}{|l|}{Load Istruction Sequence} & \multicolumn{2}{|l|}{Store Instruction Sequence} \\
\hline Fixed-point doubleword accesses & addi Idx & \[
\begin{aligned}
& \hline \mathrm{Rx}, 0, \mathrm{SI} \\
& \mathrm{Rt}, \mathrm{RA}, \mathrm{Rx}
\end{aligned}
\] & \begin{tabular}{l}
addi \\
stdx
\end{tabular} & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { RS,RA,Rx }
\end{aligned}
\] \\
\hline Floating-point as integer word accesses & \begin{tabular}{l}
addi \\
Ifiwzx
\end{tabular} & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { FRT,RA,Rx }
\end{aligned}
\] & addi stfiwx & \[
\begin{aligned}
& \hline \text { Rx,0,SI } \\
& \text { FRS,RA,Rx }
\end{aligned}
\] \\
\hline Vector byte accesses & addi Ivebx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { VRT,RA,Rx }
\end{aligned}
\] & addi stvebx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { VRS,RA,Rx }
\end{aligned}
\] \\
\hline Vector halfword accesses & addi Ivehx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { VRT,RA,Rx }
\end{aligned}
\] & addi stvehx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { VRS,RA,Rx }
\end{aligned}
\] \\
\hline Vector word accesses & addi Ivewx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { VRT,RA,Rx }
\end{aligned}
\] & addi stvewx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { VRS,RA,Rx }
\end{aligned}
\] \\
\hline Vector accesses & addi Ivx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { VRT,RA,Rx }
\end{aligned}
\] & \begin{tabular}{l}
addi \\
stvx
\end{tabular} & \[
\begin{aligned}
& \hline \text { Rx,0,SI } \\
& \text { VRS,RA,Rx }
\end{aligned}
\] \\
\hline VSX Vector accesses & addi Ixvx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XT,RA,Rx }
\end{aligned}
\] & addi stxvx & \[
\begin{aligned}
& \hline R x, 0, S I \\
& \text { XS,RA,Rx }
\end{aligned}
\] \\
\hline VSX Vector doubleword accesses & addi Ixvd2x & \[
\begin{aligned}
& \hline \mathrm{Rx}, 0, \mathrm{SI} \\
& \mathrm{XT}, \mathrm{RA}, \mathrm{Rx}
\end{aligned}
\] & addi stxvd2x & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XS,RA,Rx }
\end{aligned}
\] \\
\hline VSX Vector word accesses & addi Ixvw4x & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XT,RA,Rx }
\end{aligned}
\] & addi stxvw4x & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XS,RA,Rx }
\end{aligned}
\] \\
\hline VSX Vector halfword accesses & addi Ixvh8x & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XT,RA,Rx }
\end{aligned}
\] & addi stxvh8x & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XS,RA,Rx }
\end{aligned}
\] \\
\hline VSX Vector byte accesses & addi lxvb16x & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XT,RA,Rx }
\end{aligned}
\] & addi stxvb16x & \[
\begin{gathered}
\text { Rx,0,SI } \\
\text { xXS,RA,Rx }
\end{gathered}
\] \\
\hline VSX Vector word splat accesses & addi |xvwsx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XT,RA,Rx }
\end{aligned}
\] & n/a & \\
\hline VSX Vector doubleword splat accesses & addi Ixvdsx & \[
\begin{aligned}
& \mathrm{Rx}, 0, \mathrm{SI} \\
& \text { XT,RA,Rx }
\end{aligned}
\] & n/a & \\
\hline VSX Scalar doubleword accesses & addi Ixsdx & \[
\begin{aligned}
& \mathrm{Rx}, 0, \mathrm{SI} \\
& \mathrm{XT}, \mathrm{RA}, \mathrm{Rx}
\end{aligned}
\] & \begin{tabular}{l}
addi \\
stxsdx
\end{tabular} & \[
\begin{aligned}
& \mathrm{Rx}, 0, \mathrm{SI} \\
& \mathrm{XS}, \mathrm{RA}, \mathrm{Rx}
\end{aligned}
\] \\
\hline VSX Scalar single-precision accesses & addi lxsspx & \[
\begin{aligned}
& \mathrm{Rx}, 0, \mathrm{SI} \\
& \mathrm{XT}, \mathrm{RA}, \mathrm{Rx}
\end{aligned}
\] & addi stxsspx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XS,RA,Rx }
\end{aligned}
\] \\
\hline VSX Scalar byte accesses & addi |xsibzx & \[
\begin{aligned}
& \hline \mathrm{Rx}, 0, \mathrm{SI} \\
& \mathrm{XT}, \mathrm{RA}, \mathrm{Rx}
\end{aligned}
\] & \begin{tabular}{l}
addi \\
stxsibx
\end{tabular} & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XS,RA,Rx }
\end{aligned}
\] \\
\hline VSX Scalar halfword accesses & addi Ixsihzx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XT,RA,Rx }
\end{aligned}
\] & addi stxsihx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XS,RA,Rx }
\end{aligned}
\] \\
\hline VSX Scalar word accesses & addi Ixsiwzx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XT,RA,Rx }
\end{aligned}
\] & addi stxsiwx & \[
\begin{aligned}
& \text { Rx,0,SI } \\
& \text { XS,RA,Rx }
\end{aligned}
\] \\
\hline
\end{tabular}

Table 2: Loads and Stores with Offsets from (RA) by Magnitudes of Up to \(2^{16}\).

\section*{Programming Note}

Even independent of the performance optimization described above, the techniques illustrated in Table 1 and Table 2 generally perform better than other ways of achieving the effect of having a large displacement field for D-form and DS-form fixed-point Load/Store instructions (Table 1), and of having a displacement field for X-form and XX1-form Vector and VSX Load/Store instructions (Table 2).

The technique for the fixed-point Load/Store instructions is complicated by the fact that D-form and DS-form Loads and Stores treat the D/DS value as signed.

For simplicity, most of this Note assumes that the fixed-point Load/Store instruction is D-form; the modifications for DS-form fixed-point Load/Store instructions are straightforward.

Let the desired effective address to load from or store to be (RA) + DISP, where DISP is a signed 32-bit value.
\((\) RA \()+\) DISP \(=(\) RA \()+\) DISP \(_{0: 15}\) II DISP \(16: 31\)
\(=(R A)+\left(\right.\) DISP \(_{0: 15} \|\) Ox0000 \()+\) DISP \(_{16: 31}\)
where DISP \(_{0: 15}\) is a signed 16-bit value.
If DISP \(_{0: 15}\) is used as the SI value for the addis, the addis forms the sum
(RA) + ( DISP \(_{0: 15}\) II 0x0000)
and places the result into Rx.
If DISP \({ }_{16: 31}\) is used as the \(D\) value for the Load or Store and \(R x\) is used as the base register for the Load or Store, and DISP \(_{16}=0\), the Load or Store computes the EA to load from as
\[
\begin{aligned}
(R x)+\text { DISP }_{16: 31} & =(R A)+\left(\text { DISP }_{0: 15} I I 0 x 0000\right)+\text { DISP }_{16: 31} \\
& =(R A)+\text { DISP }^{2}
\end{aligned}
\]

However, because D-form Loads and Stores treat the D value as signed, if DISP \({ }_{16}=1\) the Load or Store computes the EA as
\[
\begin{aligned}
(R x)+\text { DISP }_{16: 31} & =(R A)+\left(\text { DISP }_{0: 15} \| 0 \times 0000\right)+\text { DISP }_{16: 31}+0 \times F F F F \_F F F F \_F F F F_{-} 0000 \\
& =(R A)+\left(\text { DISP }_{0: 15} \| 0 \times 0000\right)+\text { DISP }_{16: 31}-2^{16} \\
& =(R A)+\text { DISP }-2^{16}
\end{aligned}
\]

To compensate for this effective subtraction of \(2^{16}\), if DISP \(_{16}=1\) the SI value used for the addis must be
\(\mathrm{DISP}_{0: 15}+1\). Then the addis sets Rx to
\((\mathrm{RA})+\left(\left(\mathrm{DISP}_{0: 15}+1\right) \| 0 x 0000\right)=(\mathrm{RA})+\left(\mathrm{DISP}_{0: 15} \| 0 \times 0000\right)+2^{16}\)
and the Load or Store computes the EA as
\((R x)+\) DISP \(_{16: 31}=(R A)+\left(\right.\) DISP \(\left._{0: 15} I I 0 x 0000\right)+2^{16}+\) DISP \(_{16: 31}-2^{16}\)
\[
=(\mathrm{RA})+\text { DISP }
\]
as desired.

Thus the rules for using the technique illustrated in Table 1 are as follows.
- For the RA field of the addis, use the desired base register for the Load or Store.
- For the D field of the Load or Store, use DISP \({ }_{\text {16:31 }}\).
(For DS-form Loads and Stores, for the DS field use DISP \({ }_{16: 29}\); DISP \(_{30: 31}\) are 0b00.)
- For the SI field of the addis:
- if DISP \(_{16}=0\) use DISP \(_{0: 15}\);
- if DISP \(_{16}=1\) use DISP \(_{0: 15}+1\).

I

\subsection*{2.1.2 32-Bit Constant Generation}

The following instruction sequences will optimize performance when generating zero-extended 32-bit unsigned constants (when \(\mathrm{RA}_{0: 63}\) equal 0 ) and when performing 32-bit logical operations on \(\mathrm{RA}_{32: 63}\) ).
\begin{tabular}{|l|ll|}
\hline Operation & \multicolumn{2}{|l|}{\begin{tabular}{l} 
Instruction \\
Sequence
\end{tabular}} \\
\hline Unsigned constant & oris & \(\mathrm{Rx}, \mathrm{RA}, \mathrm{UI}_{\mathrm{h}}\) \\
(UI, UI I zero extended) & ori & \(\mathrm{Rt}, \mathrm{Rx}, \mathrm{UI}_{\mathrm{I}}\) \\
\hline Unsigned constant & xoris & \(\mathrm{Rx}, \mathrm{RA}, \mathrm{UI}_{\mathrm{h}}\) \\
(UI, UI I zero extended) & xori & \(\mathrm{Rt}, \mathrm{Rx}, \mathrm{UI}_{\mathrm{I}}\) \\
\hline
\end{tabular}

Table 3: 32-bit Unsigned Constant Generation

The following instruction sequences will optimize performance when generating 32-bit signed constants.
\begin{tabular}{|c|c|c|}
\hline Operation & \multicolumn{2}{|l|}{Instruction Sequence} \\
\hline Signed consant ( \(\mathrm{SI}_{\mathrm{h}}, \mathrm{Sl}_{\mathrm{I}}\) sign extended) & \begin{tabular}{l}
addis \\
addi
\end{tabular} & \[
\begin{aligned}
& \mathrm{Rx}, \mathrm{RA}, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{Rt}, \mathrm{Rx}, \mathrm{SI}_{\mathrm{I}}
\end{aligned}
\] \\
\hline \begin{tabular}{l}
Signed consant \\
( \(\mathrm{SI}_{\mathrm{h}}\) sign extended; Ul zero extended)
\end{tabular} & addis ori & \[
\begin{aligned}
& \mathrm{Rx}, 0, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{Rt}, \mathrm{Rx}, \mathrm{UI}_{\mathrm{I}}
\end{aligned}
\] \\
\hline
\end{tabular}

\subsection*{2.1.3 Sign and Zero Extension}

The following instruction sequences will optimize performance when converting 32-bit signed constants into 64-bit signed constants or performing other operations that require the result of an arithmetic operation to be sign extended.


The following instruction sequence will optimize performance when zero-extending the result of a 32-bit addition.
\begin{tabular}{|l|l|}
\hline Operation & \begin{tabular}{l} 
Instruction \\
Sequence
\end{tabular} \\
\hline Unsigned constant & add \\
(RA + RB zero extended) & rldicl \\
RA, Rx, 0,32 \\
\hline
\end{tabular}

Table 6: 32-bit Zero-Extended addition

\section*{Version 3.0}

\subsection*{2.1.4 Load/Store Addressing Relative to Program Counter}

The following instruction sequences will optimize performance for storage accesses to effective addresses that are offset from the CIA by magnitudes of up to \(2^{32}\).
\begin{tabular}{|c|c|c|c|c|}
\hline Operation & \multicolumn{2}{|l|}{Load Instruction Sequence} & \multicolumn{2}{|l|}{Store Instruction Sequence} \\
\hline Fixed-point byte accesses & addpcis lbz & \[
\begin{aligned}
& R x, S I_{h} \\
& R t, D(R x)
\end{aligned}
\] & addpcis stb & \[
\begin{aligned}
& \mathrm{Rx}, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{RS}, \mathrm{D}(\mathrm{Rx})
\end{aligned}
\] \\
\hline Fixed-point halfword accesses & addpcis lhz & \[
\begin{aligned}
& \mathrm{Rx}, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{Rt}, \mathrm{D}(\mathrm{Rx})
\end{aligned}
\] & addpcis sth & \[
\begin{aligned}
& R x, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{RS}, \mathrm{D}(\mathrm{Rx})
\end{aligned}
\] \\
\hline Fixed-point word accesses & addpcis Iwz & \[
\begin{aligned}
& \mathrm{Rx}, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{Rt}, \mathrm{D}(\mathrm{Rx})
\end{aligned}
\] & addpcis stw & \[
\begin{aligned}
& \mathrm{Rx}, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{RS}, \mathrm{D}(\mathrm{Rx})
\end{aligned}
\] \\
\hline Fixed-point doubleword accesses & addpcis Id & \[
\begin{aligned}
& \mathrm{Rx}, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{Rt}, \mathrm{D}(\mathrm{Rx})
\end{aligned}
\] & addpcis std & \[
\begin{aligned}
& \hline R x, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{RS}, \mathrm{D}(\mathrm{Rx})
\end{aligned}
\] \\
\hline Fixed-point doubleword accesses & addpcis Idx & \[
\begin{aligned}
& \mathrm{Rx}, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{Rt}, \mathrm{D}(\mathrm{Rx})
\end{aligned}
\] & addpcis stdx & \[
\begin{aligned}
& \hline R x, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{RS}, \mathrm{D}(\mathrm{Rx})
\end{aligned}
\] \\
\hline Floating-point single-precision accesses & addpcis Ifs & \[
\begin{aligned}
& \hline R x, S I_{h} \\
& \text { FRT,D(Rx) }
\end{aligned}
\] & addpcis stfs & \[
\begin{aligned}
& \hline R x, S I_{h} \\
& \text { FRS, } D(R x)
\end{aligned}
\] \\
\hline Floating-point double-precision accesses & addpcis Ifd & \[
\begin{aligned}
& \hline R x, S I_{h} \\
& \text { FRT,D(Rx) }
\end{aligned}
\] & addpcis stfd & \[
\begin{aligned}
& \hline R x, S I_{h} \\
& \text { FRS, } D(R x)
\end{aligned}
\] \\
\hline VSX Scalar doubleword accesses & addpcis Ixsd & \[
\begin{aligned}
& \hline \mathrm{Rx}_{\mathrm{S}, \mathrm{SI}_{\mathrm{h}}} \\
& \text { VRT,DS(Rx) }
\end{aligned}
\] & addpcis stxsd & \[
\begin{aligned}
& \text { Rx, } \mathrm{SI}_{\mathrm{h}} \\
& \text { VRS, DS(Rx) }
\end{aligned}
\] \\
\hline VSX Scalar single-precision accesses & addpcis Ixssp & \[
\begin{aligned}
& \hline \mathrm{Rx}_{\mathrm{S}, \mathrm{SI}_{\mathrm{h}}}^{\text {VRT,DS(Rx) }}
\end{aligned}
\] & addpcis stxssp & \[
\begin{aligned}
& \text { Rx, } \mathrm{SI}_{\mathrm{h}} \\
& \text { VRS, } \mathrm{DS}(\mathrm{Rx})
\end{aligned}
\] \\
\hline VSX Vector accesses & addpcis Ixv & \[
\begin{aligned}
& R x, S I_{h} \\
& \text { XT,DQ(Rx) }
\end{aligned}
\] & addpcis stxv & \[
\begin{aligned}
& \mathrm{Rx}, \mathrm{SI}_{\mathrm{h}} \\
& \mathrm{XS}, \mathrm{DQ}(\mathrm{Rx})
\end{aligned}
\] \\
\hline
\end{tabular}

Table 7: Fixed-Point, Floating-Point and VSX Load/Store Fusion with offset up to \(2^{32}\) from Program Counter
I

\section*{- Programming Note}

See the Programming Notes for Table 1.

\subsection*{2.1.5 Destructive Operation Operand Preservation}

A destructive operation is an operation that modifies one of its inputs. The VSX Vector Permute and VSX Vector Multiply-Add instructions are destructive operations because they use their destination register as a source register.
When there is a need to preserve the contents of the overwritten source register for the various VSX Vector Permute and VSX Vector Multiply-Add instructions, performance will be optimized if the xxlor instruction is used to copy the contents of the source operand into another register, and then that register is used as the destination (and source) register for the VSX Vector Permute or VSX Vector Multiply-Add instruction.

As an example, to preserve the XT source register in the xxperm instruction, the following sequence will optimize performance.
```

xxlor XC,XT,XT /* Copy (XT) to XC
xxperm XT,XA,XB /* Permute, overwriting XT

```

The set of instructions listed below, when immediately preceded by the xxlor \(\mathrm{XC}, \mathrm{XT}, \mathrm{XT}\) instruction in a sequence similar to the above example, will provide optimal performance.
\begin{tabular}{|c|c|c|}
\hline Mnemonic & & Instruction Name \\
\hline xxperm & XT,XA, XB & VSX Vector Permute \\
\hline xxpermr & XT, XA, XB & VSX Vector Permute Right Indexed \\
\hline xsmaddasp & XT, XA, XB & VSX Scalar Multiply-Add Type-A Single-Precision \\
\hline xsmsubasp & XT,XA, XB & VSX Scalar Multiply-Subtract Type-A Single-Precision \\
\hline xsnmaddasp & XT, XA, XB & VSX Scalar Negative Multiply-Add Type-A Single-Precision \\
\hline xsnmsubasp & XT, XA, XB & VSX Scalar Negative Multiply-Subtract Type-A Single-Precision \\
\hline xsmaddadp & XT, XA, XB & VSX Scalar Multiply-Add Type-A Double-Precision \\
\hline xsmsubadp & XT, XA, XB & VSX Scalar Multiply-Subtract Type-A Double-Precision \\
\hline xsnmaddadp & XT, XA, XB & VSX Scalar Negative Multiply-Add Type-A Double-Precision \\
\hline xsnmsubadp & XT, XA, XB & VSX Scalar Negative Multiply-Subtract Type-A Double-Precision \\
\hline xsmaddqp[0] & XT, XA, XB & VSX Scalar Multiply-Add Quad-Precision [using round to Odd] \\
\hline xsmsubqp[0] & XT, XA, XB & VSX Scalar Multiply-Subtract Quad-Precision [using round to Odd] \\
\hline xsnmaddqp[0] & XT,XA, XB & VSX Scalar Negative Multiply-Add Quad-Precision [using round to Odd] \\
\hline xsnmsubqp[o] & XT, XA, XB & VSX Scalar Negative Multiply-Subtract Quad-Precision [using round to Odd] \\
\hline xvmaddasp & XT, XA, XB & VSX Vector Multiply-Add Type-A Single-Precision \\
\hline xvmsubasp & XT, XA, XB & VSX Vector Multiply-Subtract Type-A Single-Precision \\
\hline xvnmaddasp & XT, XA, XB & VSX Vector Negative Multiply-Add Type-A Single-Precision \\
\hline xvnmsubasp & XT, XA, XB & VSX Vector Negative Multiply-Subtract Type-A Single-Precision \\
\hline xvmaddadp & XT, XA, XB & VSX Vector Multiply-Add Type-A Double-Precision \\
\hline xvmsubadp & XT,XA, XB & VSX Vector Multiply-Subtract Type-A Double-Precision \\
\hline xvnmaddadp & XT, XA, XB & VSX Vector Negative Multiply-Add Type-A Double-Precision \\
\hline xvnmsubadp & XT, XA, XB & VSX Vector Negative Multiply-Subtract Type-A Double-Precision \\
\hline
\end{tabular}

Table 8. VSX Multiply-Add Arithmetic Instructions Providing Optimal Performance When Preceded by xxlor

\section*{Programming Note}

Table 8 includes only the Type-A Multiply-Add instructions because supporting only one of the two types (i.e. either Type-A or Type-M) is sufficient to preserve the contents of the destination operand of the permute or Multiply-Add instruction. The xxlor instruction "preserves" the contents of the destination operand by copying it into another register, and the copy is then used as the destination operand of the Multiply-Add instruction, which is overwritten upon execution.

\subsection*{2.2 Instruction Restart}

In this section, "Load instruction" includes the Cache Management and other instructions that are stated in the instruction descriptions to be "treated as a Load", and similarly for "Store instruction".

The following instructions are never restarted after having accessed any portion of the storage operand (unless the instruction causes a "Data Address Watchpoint match", for which the corresponding rules are given in Book III).
1. A Store instruction that causes an atomic access
2. A Load instructionthat causes an atomic access to storage that is Guarded is also Caching Inhibited.

Any other Load or Store instruction may be partially executed and then aborted after having accessed a portion of the storage operand, and then re-executed (i.e., restarted, by the processor or the operating system). If an instruction is partially executed, the contents of registers are preserved to the extent that the correct result will be produced when the instruction is re-executed. Additional restrictions on the partial execution of instructions are described in Section 6.6 of Book III.

\section*{Programming Note}

In order to ensure that the contents of registers are preserved to the extent that a partially executed instruction can be re-executed correctly, the registers that are preserved must satisfy the following conditions. For any given instruction, zero or more of the conditions applies.
- For a fixed-point Load instruction that is not a multiple or string form, if \(R T=R A\) or \(R T=R B\) then the contents of register RT are not altered.
- For an update form Load or Store instruction, the contents of register RA are not altered.

\section*{Programming Note}

There are many events that might cause a Load or Store instruction to be restarted. For example, a hardware error may cause execution of the instruction to be aborted after part of the access has been performed, and the recovery operation could then cause the aborted instruction to be re-executed.

When an instruction is aborted after being partially executed, the contents of the instruction pointer indicate that the instruction has not been executed, however, the contents of some registers may have been altered and some bytes within the storage operand may have been accessed. The following are examples of an instruction being partially executed and altering the program state even though it appears that the instruction has not been executed.
1. Load Multiple, Load String: Some registers in the range of registers to be loaded may have been altered.
2. Any Store instruction, dcbz: Some bytes of the storage operand may have been altered.

\section*{Chapter 3. Management of Shared Resources}

The facilities described in this section provide the means to control the use of resources that are shared with other processors.

\subsection*{3.1 Program Priority Registers}

The Program Priority Register (PPR) is a 64-bit register that controls the program's priority. The PPR provides access to the full 64-bit PPR, and the Program Priority Register 32-bit (PPR32) provides access to the upper 32 bits of the PPR. The layouts of the PPR and PPR32 are shown in Figure 1.

PPR:
\begin{tabular}{|c|c|c|c|}
\hline //] & PRI & \multicolumn{2}{|c|}{I/I} \\
\hline 0 & 11 & 14 & 63 \\
\hline PPR32 & & & \\
\hline /I/ & PRI & III & \\
\hline 32 & 43 & \(46 \quad 63\) & \\
\hline Bit(s) & Descri & ription & \\
\hline 11:13 & \begin{tabular}{l}
Progra \\
(PPR3
\end{tabular} & ram Priority (PRI) 32 \({ }_{43: 45}\) ) & \\
\hline & 001 v & very low & \\
\hline & 010 low & low & \\
\hline & 011 m & medium low & \\
\hline & 100 m & medium & \\
\hline & 101 m & medium high & \\
\hline
\end{tabular}

Programs can always set the PRI field to very low, low, medium low, and medium priorities; programs may be allowed to set the PRI field to medium high priority during certain time intervals. (See Section 4.3.7.) If the program priority is medium high when the time interval expires or if an attempt is made to set the priority to medium high when it is not allowed, the PRI field is set to medium.

If other values are written to this field, the PRI field is not changed. (See Section 4.3.6 of Book III for additional information.)

All other fields are reserved.
Figure 1. Program Priority Register

\section*{Programming Note}

The ability to access the low-order half of the PPR (and thus the use of mfppr and mtppr) might be phased out in a future version of the architecture.

\section*{Programming Note}

By setting the PRI field, a programmer may be able to improve system throughput by causing system resources to be used more efficiently
E.g., if a program is waiting on a lock (see Section B.2), it could set low priority, with the result that more processor resources would be diverted to the program that holds the lock. This diversion of resources may enable the lock-holding program to complete the operation under the lock more quickly, and then relinquish the lock to the waiting program.

\section*{Programming Note}
or \(R x, R x, R x\) can be used to modify the PRI field; see Section 3.2.

\section*{Programming Note}

When the system error handler is invoked, the PRI field may be set to an undefined value.

\section*{3.2 "or" Instruction}

\section*{Setting the PPR}

The or \(R x, R x, R x\) (see Book I) instruction can be used to set PPR \(_{\text {PRI }}\) as shown in Figure . or. \(R x, R x, R x\) does not set PPR PRII .
\begin{tabular}{|c|c|l|}
\hline \(\mathbf{R x}\) & \(\mathbf{P P R}_{\mathbf{P R I}}\) & Priority \\
\hline 31 & 001 & very low \\
\hline 1 & 010 & low \\
\hline 6 & 011 & medium low \\
\hline 2 & 100 & medium \\
\hline 5 & 101 & medium high \\
\hline
\end{tabular}

\section*{Version 3.0}

\section*{Priority levels for or \(R x, R x, R x\)}

Programs can always set the PRI field to very low, low, medium low, and medium priorities; programs may be allowed to set the PRI field to medium high priority during certain time intervals. (See Section 4.3.7 of Book III.) If the program priority is medium high when the time interval expires or if an attempt is made to set the priority to medium high when it is not allowed, the PRI field is set to medium.

\section*{Programming Note}

Warning: Other forms of or \(R x, R x, R x\) that are not described in this section and in Section 4.3.3 may also cause program priority to change. Use of these forms should be avoided except when software explicitly intends to alter program priority. If a no-op is needed, the preferred no-op (ori \(0,0,0\) ) should be used.

\title{
Chapter 4. Storage Control Instructions
}

\subsection*{4.1 Parameters Useful to Application Programs}

It is suggested that the operating system provide a service that allows an application program to obtain the following information.
1. The virtual page sizes
2. Coherence block size
3. Reservation granule size
4. An indication of the cache model implemented (e.g., Harvard-style cache, combined cache)
5. Instruction cache size
6. Data cache size
7. Instruction cache block size
8. Data cache block size
9. Instruction cache associativity
10. Data cache associativity
11. Number of stream IDs supported for the stream variant of dcbt
12. Factors for converting the Time Base to seconds
13. Maximum transaction level

If the caches are combined, the same value should be given for an instruction cache attribute and the corresponding data cache attribute.

\subsection*{4.2 Data Stream Control Register (DSCR)}

The layout of the Data Stream Control Register (DSCR) is shown in Figure 2 below.


Figure 2. Data Stream Control Register
Bit(s) Description
39 Software Transient Enable (SWTE)
0 SWTE is disabled.

1 Applies the transient attribute to soft-ware-defined streams.

Hardware Transient Enable (HWTE)
0 HWTE is disabled.
1 Applies the transient attribute to hard-ware-detected streams.

\section*{Store Transient Enable (STE)}

0 STE is disabled.
1 Applies the transient attribute to store streams.

\section*{Load Transient Enable (LTE)}

0 LTE is disabled.
1 Applies the transient attribute to load streams.

\section*{Software Unit count Enable (SWUE)}

0 SWUE is disabled.
1 Applies the unit count to software-defined streams.

\section*{Hardware Unit count Enable (HWUE)}

0 HWUE is disabled.
1 Applies the unit count to hard-ware-detected streams.

Depth Attainment Urgency (URG)
This field indicates how quickly the prefetch depth should be reached for hard-ware-detected streams. Values and their meanings are as follows.

0 default
1 not urgent
2 least urgent
3 less urgent
4 medium
5 urgent
6 more urgent
7 most urgent
Load Stream Disable (LSD)
0 No effect.

1 Disables hardware detection and initiation of load streams.

\section*{Stride-N Stream Enable (SNSE)}

0 No effect.
1 Enables the hardware detection and initiation of load and store streams that have a stride greater than a single cache block. Such load streams are detected only when LSD is also zero. Such store streams are detected only when SSE is also one.

\section*{Store Stream Enable (SSE)}

0 No effect.
1 Enables hardware detection and initiation of store streams.

\section*{61:63 Default Prefetch Depth (DPFD)}

This field supplies a prefetch depth for hard-ware-detected streams and for soft-ware-defined streams for which a depth of zero is specified or for which dcbt/dcbtst with TH=1010 is not used in their description. Values and their meanings are as follows.

0 default ( LPCR \(_{\text {DPFD }}\) )
1 none
2 shallowest
3 shallow
4 medium
5 deep
6 deeper
7 deepest
The contents of the DSCR affect how a processor handles hardware-detected and software-defined data streams. The DSCR provides the only means by which software can control or supply information for hard-ware-detected data streams. The DPFD, UNITCNT, and transient fields may also be used instead of the \(\mathrm{TH}=01010\) variant of dcbt for software-defined data streams, especially when multiple streams have these attributes in common. See Section 4.3.2, "Data Cache Instructions" on page 843, for information on streams and how software may specify them.

\section*{Programming Note}

The URG, LSD, SNSE and SSE fields do not affect the initiation of streams specified using the dcbt and dcbtst instructions.

Note that even when SNSE is not set, hardware may detect Stride-N streams in intervals when they access elements that map to sequential cache blocks.

> - Programming Note
> In order for the DSCR to apply the transient attribute to streams, at least two of the four enable bits must be set: one to choose a type of access (load or store), and one to choose a kind of prefetching (software-defined or hardware-detected).

\section*{Programming Note}

The purpose of Depth Attainment Urgency is to regulate the rate of prefetch generation from the cycle at which the hardware first detects an incipient stream until the cycle when the prefetch Depth is reached. A more urgent setting will benefit applications that are dominated by short to medium length streams, because otherwise prefetching does not occur rapidly enough to benefit them. In contrast, applications that frequently cause unproductive prefetches due to stream mispredicts will benefit from a less urgent setting.
Unlike the Depth, the Depth Attainment Urgency applies only to hardware-detected streams. Furthermore, the DSCR provides the only point of control for this parameter. Software-defined streams are assumed not to have the correctness risk associated with hardware streams, and therefore are set to reach their depth relatively quickly.

\section*{Programming Note}

In versions of the architecture that precede Version 2.07, mtspr specifying the DSCR caused all active and nascent data streams to cease to exist. In those versions of the architecture, the DSCR was used as an overall control mechanism to specify a single global profile for all streams. Beginning with Version 2.07, the DSCR is intended to control and accelerate the creation of new streams without disturbing existing streams.

\subsection*{4.3 Cache Management Instructions}

The Cache Management instructions obey the sequential execution model except as described in Section 4.3.1.

In the instruction descriptions the statements "this instruction is treated as a Load" and "this instruction is treated as a Store" mean that the instruction is treated as a Load (Store) from (to) the addressed byte with respect to address translation, the definition of program order on page 811, storage protection, reference and change recording, and Performance Monitor events (see Section 9.4.5 of Book III).

\section*{Programming Note}

Accesses that are caused by or associated with Cache Management instructions that are "treated as a Load" or "treated as a Store" are not subject to the special ordering rules described for SAO storage. These accesses are always performed in accordance with the weakly consistent storage model.

Some Cache Management instructions contain a CT field that is used to specify a cache level within a cache hierarchy or a portion of a cache structure to which the instruction is to be applied. The correspondence between the CT value specified and the cache level is shown below.
\begin{tabular}{ll} 
CT Field Value & Cache Level \\
0 & Primary Cache \\
2 & Secondary Cache
\end{tabular}

CT values not shown above may be used to specify implementation-dependent cache levels or implemen-tation-dependent portions of a cache structure.

\subsection*{4.3.1 Instruction Cache Instructions}

\section*{Instruction Cache Block Invalidate X-form}
icbi RA,RB
\begin{tabular}{|l|l|l|c|c|c|l|}
\hline 31 & \multicolumn{1}{|c|}{ I/I } & RA & RB & & 982 & 1 \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}

Let the effective address (EA) be the sum (RAIO)+(RB).
If the block containing the byte addressed by EA is in storage that is Memory Coherence Required and a block containing the byte addressed by EA is in the instruction cache of any processors, the block is invalidated in those instruction caches.

If the block containing the byte addressed by EA is in storage that is not Memory Coherence Required and the block is in the instruction cache of this processor, the block is invalidated in that instruction cache.
The function of this instruction is independent of whether the block containing the byte addressed by EA is in storage that is Write Through Required or Caching Inhibited.

This instruction is treated as a Load (see Section 4.3), except that reference and change recording need not be done.

\section*{Special Registers Altered:}

None

\section*{Programming Note}

Because the instruction is treated as a Load, the effective address is translated using translation resources that are used for data accesses, even though the block being invalidated was copied into the instruction cache based on translation resources used for instruction fetches (see Book III).

\section*{Programming Note}

The invalidation of the specified block need not have been performed with respect to the processor executing the icbi instruction until a subsequent isync instruction has been executed by that processor. No other instruction or event has the corresponding effect.

\section*{Instruction Cache Block Touch}

X-form
icbt \(\quad \mathrm{CT}, \mathrm{RA}, \mathrm{RB}\)
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline 31 & 1 & CT & RA & \multicolumn{1}{c|}{ RB } & & 22 & 1 \\
0 & & 7 & 7 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}

Let the effective address (EA) be the sum (RAIO)+(RB).
The icbt instruction provides a hint that the program will probably soon execute code from the block containing the byte addressed by EA, and that the block containing the byte addressed by EA is to be loaded into the cache specified by the CT field. (See Section 4.3 of Book II.) If the CT field is set to a value not supported by the implementation, no operation is performed.

The hint is ignored if the block is Caching Inhibited.
This instruction treated as a Load (see Section 4.3), except that the system data storage error handler is not invoked, and reference and change recording need not be done.

\section*{Special Registers Altered:}

None

\subsection*{4.3.2 Data Cache Instructions}

The Data Cache instructions control various aspects of the data cache.

\section*{TH field in the dcbt and dcbtst instructions}

Described below are the TH field values for the dcbt and dcbtst instructions. For all TH field values which are not listed, the hint provided by the instruction is undefined.

\section*{TH=0b00000}

If \(\mathrm{TH}=0 \mathrm{~b} 00000\), the \(\boldsymbol{d c b t} / \boldsymbol{d c b t s t}\) instruction provides a hint that the program will probably soon access the block containing the byte addressed by EA.

\section*{TH=0b01000 - Ob01111}

The dcbt/dcbtst instructions provide hints regarding a sequence of accesses to data elements, or indicate the expected use thereof. Such a sequence is called a "data stream", and a dcbt/dcbtst instruction in which TH is set to one of these values is said to be a "data stream variant" of dcbt/dcbtst. In the remainder of this section, "data stream" may be abbreviated to "stream".

A data stream to which a program may perform Load accesses is said to be a "load data stream", and is described using the data stream variants of the dcbt instruction. A data stream to which a program may perform Store accesses is said to be a "store data stream", and is described using the data stream variants of the dcbtst instruction.

When, and how often, effective addresses for a data stream are translated is implementation-dependent.
Each data element is associated with a unit of storage, which is the aligned 128-byte location in storage that contains the first byte of the element. The data stream variants may be used to specify the address of the beginning of the data stream, the displacement (stride) between the first byte of successive elements, and the number of unique units of storage that are associated with all of the data elements. If the stride is specified, both the stride and the address of the first element are specified at 4 byte granularity. If the stride is not specified, the address of the first element is the address of the first unit.

\section*{Programming Note}

The architecture does not provide a way to specify the size of the data elements that compose a stream. An implementation may assume some fixed size for all data elements. As a result, depending on the offset, stride, and size (and in particular whether the elements are aligned), the implementation may reduce the latency for accessing only a portion of some of the elements. A future version of the architecture may enable the specification of element size to avoid this limitation.

Each such data stream is associated, by software, with a stream ID, which is a resource that the processor uses to distinguish the data stream from other such data streams. The number of stream IDs is an imple-mentation-dependent value in the range \(1: 16\). Stream IDs are numbered sequentially starting from 0 .

The encodings of the TH field and of the corresponding EA values are as follows. In the EA layout diagrams, fields shown as "/"s are reserved. These reserved fields are treated in the same manner as the corresponding case for instruction fields (see Section 1.3.3 of Book I). If a reserved value is specified for a defined EA field, or if a TH value is specified that is not explicitly defined below, the hint provided by the instruction is undefined.

\section*{TH Description}

01000 The dcbt/dcbtst instruction provides a hint that describes certain attributes of a data stream, and may indicate that the program will probably soon access the stream.

The EA is interpreted as follows.
\begin{tabular}{|l|l|l|l|}
\hline & EATRUNC & D UG & ID \\
\hline 0 & \(57 \quad 596063\)
\end{tabular}

\section*{Bit(s) Description}

0:56 EATRUNC
High-order 57 bits of the effective address of the first element of the data stream. (i.e., the effective address of the first unit of the stream is EATRUNC \(\|{ }^{7} 0\) )
57 Direction (D)
0 Subsequent elements have increasing addresses.
1 Subsequent elements have decreasing addresses.

58 Unlimited/GO (UG)
0 No information is provided by the UG field.
1 The number of elements in the data stream is unlimited, the elements are adjacent to each other, the program's need for each element of the stream is not likely to be transient, and the program will probably soon access the stream.
59 Reserved
60:63 Stream ID (ID)
Stream ID to use for this data stream.
01010 The dcbt/dcbtst instruction provides a hint that describes certain attributes of a data stream, or indicates that the program will probably soon access data streams that have been described using data stream variants of the dcbt/dcbtst instruction, or will probably no longer access such data streams.

The EA is interpreted as follows. If GO=1 and \(\mathrm{S} \neq 0 \mathrm{~b} 00\) the hint provided by the instruction is undefined; the remainder of this instruction description assumes that this combination is not used.


Bit(s) Description
0:31 Reserved
32 GO
0 No information is provided by the GO field.
1 For dcbt, the program will probably soon access all nascent load and store data streams that have been completely described, and will probably no longer access all other nascent load and store data streams. All other fields of the EA are ignored. ("Nascent" and "completely described" are defined below.) For dcbtst, this field value holds no meaning and is treated as though it were zero.

\section*{33:34 Stop (S)}

00 No information is provided by the \(S\) field.
01 Reserved
10 The program will probably no longer access the data stream (if any) associated with the specified
stream ID. (All other fields of the EA except the ID field are ignored.)
11 For dcbt, the program will probably no longer access the load and store data streams associated with all stream IDs. (All other fields of the EA are ignored.) For dcbtst, this field value holds no meaning, and is treated as though it were ObOO.

\section*{35 Reserved}

36:38 Depth (DEP)
The DEP field provides a relative estimate of how many elements ahead of the point of stream use the latency-reducing actions should go. This value reflects a comparison of the rate of consumption of the elements of the data stream and the latency to bring an arbitrary element of the stream into cache. The values are as follows.
\begin{tabular}{ll}
0 & default = DSCR \\
1 & none \\
2 & shallowest \\
3 & shallow \\
4 & medium \\
5 & deep \\
6 & deeper \\
7 & deepest
\end{tabular}

39:46 Reserved

\section*{47:56 UNITCNT}

Number of units in data stream.
\(57 \quad\) Transient ( T )
If \(T=1\), the program's need for each element of the data stream is likely to be transient (i.e., the time interval during which the program accesses the element is likely to be short).
58 Unlimited (U)
If \(\mathrm{U}=1\), the number of units in the data stream is unlimited (and the UNITCNT field is ignored).
59 Reserved
60:63 Stream ID (ID)
Stream ID to use for this data stream ( \(\mathrm{GO}=0\) and \(\mathrm{S}=0 \mathrm{~b} 00\) ), or stream ID associated with the data stream which the program will probably no longer access(S=0b10).

\section*{Programming Note}

To maximize the utility of the Depth control mechanism, the architecture provides a hierarchy of three ways to program it. The DPFD field in the LPCR is used by the provisory/firmware to set a safe or appropriate default depth for unaware operating systems and applications. The DPFD field in the DSCR may be initialized by the aware OS and overwritten by an application via the OS-provided service when per stream control is unnecessary or unaffordable. The DEP field in the EA specification when \(\mathrm{TH}=0 \mathrm{~b} 01010\) may be used by the application to specify the depth on a per-stream basis.

The number of elements ahead of the point of stream use indicated by a given depth value may differ across implementations, as may the latency to bring a given element into the cache. To achieve optimum performance, some experimentation with different depth values may be necessary.

01011 The dcbt/dcbtst instruction provides a hint that describes certain attributes of a data stream.

The EA is interpreted as follows.
\begin{tabular}{|l|l|l|l|l|}
\hline\(/ / /\) & STRIDE & OFFSET & // & \multicolumn{1}{|c|}{ ID } \\
\hline 0 & 32 & 50 & 56 & 60
\end{tabular}

\section*{Bit(s) Description}

0:31 Reserved
32:49 Stride
The displacement, in words, between the first byte of successive elements in the stream. The effective address of the \(\mathrm{N}^{\text {th }}\) element in the stream is
( \(\mathrm{N}-1\) ) \(\times\) STRIDE \(\times 4\)
greater than or less than the effective address of the first element of the stream, depending on the direction specified for the stream.

50 Reserved
51:55 Offset
The word-offset of the first element of the stream in its unit (i.e., the effective address of the first element of the stream is (EATRUNC II OFFSET II 0b00)).

56:59Reserved
60:63 Stream ID (ID)
Stream ID to use for this data stream.

\section*{Programming Note}

A program should use a dcbt/dcbtst instruction with \(\mathrm{TH}=0 \mathrm{~b} 01011\) only when the stride is larger than 128 bytes. Otherwise, consecutive units will be accessed, so the additional stream information has no benefit.

If the specified stream ID value is greater than \(m-1\), where \(m\) is the number of stream IDs provided by the implementation, and either (a) \(\mathrm{TH}=0 \mathrm{~b} 01000\) or \(\mathrm{TH}=0 \mathrm{~b} 01011\), or (b) \(\mathrm{TH}=0 \mathrm{~b} 01010\) with \(\mathrm{GO}=0\) and \(\mathrm{S} \neq 0 \mathrm{~b} 11\), no hint is provided by the instruction.

The following terminology is used to describe the state of a data stream. Except as described in the paragraph after the next paragraph, the state of a data stream at a given time is determined by the most recently provided hint(s) for the stream.
- A data stream for which only descriptive hints have been provided (by dcbt/dcbtst instructions with \(\mathrm{TH}=0 \mathrm{~b} 01000\) and \(\mathrm{UG}=0, \mathrm{TH}=0 \mathrm{~b} 01010\) and \(\mathrm{GO}=0\) and \(\mathrm{S}=0 \mathrm{~b} 00\), and/or with \(\mathrm{TH}=0 \mathrm{bO1011}\) ) is said to be "nascent". A nascent data stream for which all relevant descriptive hints have been provided (by the dcbt/dcbtst usages listed in the preceding sentence) is considered to be "completely described". The order of descriptive hints with respect to one another is unimportant.
- A data stream for which a hint has been provided (by a dcbt/dcbtst instruction with \(\mathrm{TH}=0 \mathrm{~b} 01000\) and UG=1 or dcbt with \(\mathrm{TH}=0 \mathrm{b01010}\) and GO=1) that the program will probably soon access it is said to be "active".
- A data stream that is either nascent or active is considered to "exist".
- A data stream for which a hint has been provided (e.g., by a dcbt instruction with TH=0b01010 and \(\mathrm{S} \neq 0 \mathrm{b00}\) ) that the program will probably no longer access it is considered no longer to exist.
The hint provided by a dcbt/dcbtst instruction with TH=0b01000 and UG=1 implicitly includes a hint that the program will probably no longer access the data stream (if any) previously associated with the specified stream ID. The hint provided by a dcbt/dcbtst instruction with \(\mathrm{TH}=0 \mathrm{~b} 01000\) and \(\mathrm{UG}=0\), or with \(\mathrm{TH}=0 \mathrm{~b} 01010\) and \(\mathrm{GO}=0\) and \(\mathrm{S}=0 \mathrm{~b} 00\), or with \(\mathrm{TH}=0 \mathrm{~b} 01011\) implicitly includes a hint that the program will probably no longer access the active data stream (if any) previously associated with the specified stream ID.
If a data stream is specified without using a dcbt/ dcbtst instruction with TH=0b01010 and GO=0 and \(\mathrm{S}=0 \mathrm{~b} 00\), then the number of elements in the stream is unlimited, and the program's need for each element of the stream is not likely to be transient. If a data stream is specified without using a dcbt/dcbtst instruction with

\section*{Version 3.0}

TH=0b01011, then the stream will access consecutive units of storage.
Interrupts (see Book III) cause all existing data streams to cease to exist. In addition, depending on the implementation, certain conditions and events may cause an existing data stream to cease to exist; for example, in some implementations an existing data stream ceases to exist when it comes to the end of a page.

\section*{Programming Note}

To obtain the best performance across the widest range of implementations that support the data stream variants of dcbt/dcbtst, the programmer should assume the following model when using those variants.
- The processor's response to a hint that the program will probably soon access a given data stream is to take actions that reduce the latency of accesses to the first few elements of the stream. (Such actions may include prefetching cache blocks into levels of the storage hierarchy that are "near" the processor.) Thereafter, as the program accesses each successive element of the stream, the processor takes latency-reducing actions for additional elements of the stream, pacing these actions with the program's accesses (i.e., taking the actions for only a limited number of elements ahead of the element that the program is currently accessing).

The processor's response to a hint that the program will probably no longer access a given data stream, or to the cessation of existence of a data stream, is to stop taking latency-reducing actions for the stream.
- A data stream having finite length ceases to exist when the latency-reducing actions have been taken for all elements of the stream.
- If the program ceases to need a given data stream before having accessed all elements of the stream (always the case for streams having unlimited length), performance may be improved if the program then provides a hint that it will no longer access the stream (e.g., by executing the appropriate dcbt instruction with \(\mathrm{TH}=0 \mathrm{~b} 01010\) and \(\mathrm{S} \neq 0 \mathrm{~b} 00\) ).
- At each level of the storage hierarchy that is "near" the processor, elements of a data stream that is specified as transient are most likely to be replaced. As a result, it may be desirable to stagger addresses of streams (choose addresses that map to different cache congruence classes) to reduce the likelihood that an element of a transient stream will be replaced prior to being accessed by the program.
- Processors that comply with versions of the architecture that do not support the TH field at all treat \(\mathrm{TH}=0 \mathrm{~b} 01000,0 \mathrm{~b} 01010\), and \(0 \mathrm{b01011}\) as if \(\mathrm{TH}=\) Ob00000.
- A single set of stream IDs is shared between the \(\boldsymbol{d c b t}\) and dcbtst instructions.
- On some implementations, data streams that are not specified by software may be detected by the processor. Such data streams are called "hard-ware-detected data streams". On some such implementations, data stream resources (resources that are used primarily to support data streams) are shared between software-specified data streams and hardware-detected data streams. On these latter implementations, the programming model includes the following.
- Software-specified data streams take precedence over hardware-detected data streams in use of data stream resources.
- The processor's response to a hint that the program will probably no longer access a given data stream, or to the cessation of existence of a data stream, includes releasing the associated data stream resources, so that they can be used by hardware-detected data streams.

\section*{Version 3.0}

\section*{Programming Note}

The latency-reducing actions taken in response to a program's hints about access to a data stream, including the depth and urgency parameters, may vary based on its behavior and on the behavior of other programs sharing platform resources, as well as on the design of the platform resources they use. Without actually changing the stream specification or DSCR parameters, the processor may adjust its actions (e.g. slow down prefetches or be more selective choosing them) based on their effectiveness and on the availability of storage bandwidth. In general, the goal of this variation is to improve overall system performance and fairness across the set of programs that share resources. There often will be a performance benefit, however, from adjusting stream specifications to the platform and co-resident programs to adjust for these actions by the processor.

\section*{Programming Note}

This Programming Note describes several aspects of using the data stream variants of the dcbt and dcbtst instructions.
- A non-transient data stream having unlimited length and which will access consecutive units in storage can be completely specified, including providing the hint that the program will probably soon access it, using one dcbt instruction. The corresponding specification for a data stream having other attributes requires two or three dcbt/dcbtst instructions to describe the stream and one additional dcbt instruction to start the stream. However, one dcbt instruction with TH=0b01010 and \(\mathrm{GO}=1\) can apply to a set of the data streams described in the preceding sentence, so the corresponding specification for n such data streams requires \(2 \times n\) to \(3 \times n\) dcbt/dcbtst instructions plus one dcbt instruction. (There is no need to execute a dcbt/dcbtst instruction with TH=0b01010 and \(\mathrm{S}=0 \mathrm{~b} 10\) for a given stream ID before using the stream ID for a new data stream; the implicit portion of the hint provided by dcbt/dcbtst instructions that describe data streams suffices.)
■ If it is desired that the hint provided by a given dcbt/dcbtst instruction be provided in program order with respect to the hint provided by another dcbt/dcbtst instruction, the two instructions must be separated by an eieio instruction. For example, if a dcbt instruction with \(\mathrm{TH}=0 \mathrm{~b} 01010\) and \(\mathrm{GO}=1\) is intended to indicate that the program will probably soon access nascent data streams described (completely) by preceding dcbt/dcbtst instructions, and is intended not to indicate that the program will probably soon access nascent data streams described (completely) by following dcbt/ dcbtst instructions, an eieio instruction must sep-
arate the \(d c b t\) instruction with \(\mathrm{GO}=1\) from the preceding dcbt/dcbtst instructions, and another eieio instruction must separate that dcbt instruction from the following dcbt/dcbtst instructions.
■ In practice, the second eieio described above can sometimes be omitted. For example, if the program consists of an outer loop that contains the dcbt/ dcbtst instructions and an inner loop that contains the Load or Store instructions that access the data streams, the characteristics of the inner loop and of the implementation's branch prediction mechanisms may make it highly unlikely that hints corresponding to a given iteration of the outer loop will be provided out of program order with respect to hints corresponding to the previous iteration of the outer loop. (Also, any providing of hints out of program order affects only performance, not program correctness.)
■ To mitigate the effects of interrupts on data streams, it may be desirable to specify a given "logical" data stream as a sequence of shorter, component data streams. Similar considerations apply to conditions and events that, depending on the implementation, may cause an existing data stream to cease to exist; for example, in some implementations an existing data stream ceases to exist when it comes to the end of a virtual page.
- If it is desired to specify data streams without regard to the number of stream IDs provided by the implementation, stream IDs should be assigned to data streams in order of decreasing stream importance (stream ID 0 to the most important stream, stream ID 1 to the next most important stream, etc.). This order ensures that the hints for the most important data streams will be provided.

\section*{TH=Ob10000}

If \(\mathrm{TH}=0 \mathrm{~b} 10000\), the \(\boldsymbol{d} \boldsymbol{c b} \boldsymbol{t}\) instruction provides a hint that the program will probably soon load from the block containing the byte addressed by EA, and that the program's need for the block will be transient (i.e., the time interval during which the program accesses the block is likely to be short).

\section*{Programming Note}

The processor's response to the hint that access to the block will be transient is to prefetch data into the cache hierarchy in a way that minimizes the displacement of data that has not been identified as transient.

\section*{TH=0b10001}

If \(\mathrm{TH}=0 \mathrm{~b} 10001\), the \(\boldsymbol{d c b t}\) instruction provides a hint that the program will probably not access the block containing the byte addressed by EA for a relatively long period of time.

Version 3.0
Data Cache Block Touch X-form
dcbt RA,RB,TH
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & \multicolumn{1}{|c|}{ TH } & \multicolumn{1}{|c|}{ RA } & RB & \multicolumn{2}{|c|}{278} & \(/\) \\
0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}

Let the effective address (EA) be the sum (RAIO)+(RB).
The dcbt instruction provides a hint that describes a block or data stream to which the program may perform a Load access. The instruction is also used to indicate imminent access or end of access to described load and store data streams. A hint that the program will probably soon load from a given storage location is ignored if the location is Caching Inhibited or Guarded.
The only operation that is "caused" by the dcbt instruction is the providing of the hint. The actions (if any) taken by the processor in response to the hint are not considered to be "caused by" or "associated with" the \(\boldsymbol{d} \boldsymbol{c} \boldsymbol{b} \boldsymbol{t}\) instruction (e.g., dcbt is considered not to cause any data accesses). No means are provided by which software can synchronize these actions with the execution of the instruction stream. For example, these actions are not ordered by the memory barrier created by a sync instruction.

The dcbt instruction may complete before the operation it causes has been performed.
The nature of the hint depends, in part, on the value of the TH field, as specified at the beginning of this section. If \(\mathrm{TH} \neq 0 \mathrm{~b} 01010\) and \(\mathrm{TH} \neq 0 \mathrm{~b} 01011\), this instruction is treated as a Load (see Section 4.3), except that the system data storage error handler is not invoked, and reference and change recording need not be done.

\section*{Special Registers Altered:}

None

\section*{Extended Mnemonics:}

Extended mnemonics are provided for the Data Cache Block Touch instruction so that it can be coded with the TH value as the last operand for all categories, and so that the transient hint can be specified without coding the TH field explicitly.
\(\left.\begin{array}{ll}\text { Extended: } & \text { Equivalent to: } \\ \text { dcbtct RA,RB,TH } & \begin{array}{l}\text { dcbt for TH values of 0b00000 - } \\ \text { Ob00111; }\end{array} \\ & \begin{array}{l}\text { other TH values are invalid. }\end{array} \\ \text { dcbtds RA,RB,TH } & \begin{array}{l}\text { dcbt for TH values of Ob00000 or } \\ \text { Ob01000 - Ob01111; }\end{array} \\ \text { other TH values are invalid. }\end{array}\right\}\)

\section*{Programming Notes}

New programs should avoid using the dcbt and dcbtst mnemonics; one of the extended mnemonics should be used exclusively.
If the dcbt mnemonic is used with only two operands, the TH operand is assumed to be 0b00000.

Processors that comply with versions of the architecture that precede Version 2.01 do not necessarily ignore the hint provided by dcbt and dcbtst if the specified block is in storage that is Guarded and not Caching Inhibited.

\section*{Programming Note}

See the Programming Notes at the beginning of this section.

\section*{Data Cache Block Touch for Store X-form}

> dcbtst RA,RB,TH
\begin{tabular}{|l|l|l|l|c|c|l|}
\hline 31 & TH & RA & RB & & 246 & \(/\) \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}

Let the effective address (EA) be the sum (RAIO)+(RB).
The dcbtst instruction provides a hint that describes a block or data stream to which the program may perform a Store access, or indicates the expected use thereof. A hint that the program will soon store to a given storage location is ignored if the location is Caching Inhibited or Guarded.

The only operation that is "caused by" the dcbtst instruction is the providing of the hint. The actions (if any) taken by the processor in response to the hint are not considered to be "caused by" or "associated with" the dcbtst instruction (e.g., dcbtst is considered not to cause any data accesses). No means are provided by which software can synchronize these actions with the execution of the instruction stream. For example, these actions are not ordered by memory barriers.
The dcbtst instruction may complete before the operation it causes has been performed.

The nature of the hint depends, in part, on the value of the TH field, as specified at the beginning of this section. If \(\mathrm{TH} \neq 0 \mathrm{~b} 01010\) and \(\mathrm{TH} \neq 0 \mathrm{~b} 01011\), this instruction is treated as a Store (see Section 4.3), except that the system data storage error handler is not invoked, reference recording need not be done, and change recording is not done.

\section*{Special Registers Altered:}

None

\section*{Extended Mnemonics:}

Extended mnemonics are provided for the Data Cache Block Touch for Store instruction so that it can be coded with the TH value as the last operand for all categories, and so that the transient hint can be specified without coding the TH field explicitly.
\begin{tabular}{lc} 
Extended: & Equivalent to: \\
dcbtstct RA,RB,TH & \begin{tabular}{c} 
dcbtst for TH values of 0b00000 \\
or Ob00000 - Ob00111; \\
other TH values are invalid.
\end{tabular} \\
dcbtstds RA,RB,TH & \begin{tabular}{c} 
dcbtst for TH values of 0b00000 \\
or Ob01000 - Ob01111;
\end{tabular} \\
other TH values are invalid.
\end{tabular}
extended:
dcbtstct RA,RB,TH
dcbtstds RA,RB,TH
dcbtstt RA,RB
dcbtst for TH value of Ob10000.

\section*{Programming Note}

See the Programming Notes at the beginning of this section.

\section*{Data Cache Block set to Zero}

X-form
dcbz RA,RB
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & \multicolumn{1}{|c|}{ III } & RA & RB & & 1014 & 1 \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

if $R A=0$ then $b \leftarrow 0$
else $\quad \mathrm{b} \leftarrow(\mathrm{RA})$
$\mathrm{EA} \leftarrow \mathrm{b}+(\mathrm{RB})$
$\mathrm{n} \leftarrow$ block size (bytes)
$\mathrm{m} \leftarrow \log _{2}(\mathrm{n})$
$\mathrm{ea} \leftarrow E A_{0: 63-\mathrm{m}_{\mathrm{n}}}\| \|^{\mathrm{m}} 0$
$\operatorname{MEM}($ ea, n$) \stackrel{\text { n }}{\leftarrow} \mathrm{n}_{0 \times 00}$

```

Let the effective address (EA) be the sum (RAIO)+(RB).
All bytes in the block containing the byte addressed by EA are set to zero.

This instruction is treated as a Store (see Section 4.3).

\section*{Special Registers Altered: \\ None}

\section*{Programming Note}
dcbz does not cause the block to exist in the data cache if the block is in storage that is Caching Inhibited.

For storage that is neither Write Through Required nor Caching Inhibited, dcbz provides an efficient means of setting blocks of storage to zero. It can be used to initialize large areas of such storage, in a manner that is likely to consume less memory bandwidth than an equivalent sequence of Store instructions.

For storage that is either Write Through Required or Caching Inhibited, dcbz is likely to take significantly longer to execute than an equivalent sequence of Store instructions. For example, on some implementations dcbz for such storage may cause the system alignment error handler to be invoked; on such implementations the system alignment error handler sets the specified block to zero using Store instructions.

See Section 5.9.1 of Book III for additional information about dcbz.

Data Cache Block Store
X-form
dcbst RA,RB
\begin{tabular}{|l|l|l|l|l|ll|l|}
\hline & 31 & & I/I & RA & RB & & 54 \\
\hline 0 & & 6 & 11 & 16 & 21 & & 31 \\
\hline
\end{tabular}

Let the effective address (EA) be the sum (RAIO)+(RB).
If the block containing the byte addressed by EA is in storage that is Memory Coherence Required and a block containing the byte addressed by EA is in the data cache of any processor and any locations in the block are considered to be modified there, those locations are written to main storage, additional locations in the block may be written to main storage, and the block ceases to be considered to be modified in that data cache.

If the block containing the byte addressed by EA is in storage that is not Memory Coherence Required and the block is in the data cache of this processor and any locations in the block are considered to be modified there, those locations are written to main storage, additional locations in the block may be written to main storage, and the block ceases to be considered to be modified in that data cache.

The function of this instruction is independent of whether the block containing the byte addressed by EA is in storage that is Write Through Required or Caching Inhibited.

This instruction is treated as a Load (see Section 4.3), except that reference and change recording need not be done.

\section*{Special Registers Altered:}

None

\section*{Data Cache Block Flush}

X-form
dcbf RA,RB,L
\begin{tabular}{|l|l|l|l|l|l|ll|l|}
\hline 31 & I/I & L & RA & RB & & 86 & \\
\hline 0 & & 6 & 9 & 11 & 16 & & 21 & \\
31 \\
\hline
\end{tabular}

Let the effective address (EA) be the sum (RAIO)+(RB). L=0

If the block containing the byte addressed by EA is in storage that is Memory Coherence Required and a block containing the byte addressed by EA is in the data cache of any processor and any locations in the block are considered to be modified there, those locations are written to main storage and additional locations in the block may be written to main storage. The block is invalidated in the data caches of all processors.

If the block containing the byte addressed by EA is in storage that is not Memory Coherence Required and the block is in the data cache of this processor and any locations in the block are considered to be modified there, those locations are written to main storage and additional locations in the block may be written to main storage. The block is invalidated in the data cache of this processor.

\section*{L=1 ("dcbf local")}

The \(L=1\) form of the dcbf instruction permits a program to limit the scope of the "flush" operation to the data cache of this processor. If the block containing the byte addressed by EA is in the data cache of this processor, it is removed from this cache. The coherence of the block is maintained to the extent required by the Memory Coherence Required storage attribute.
L = 3 ("dcbf local primary")
The \(L=3\) form of the dcbf instruction permits a program to limit the scope of the "flush" operation to the primary data cache of this processor. If the block containing the byte addressed by EA is in the primary data cache of this processor, it is removed from this cache. The coherence of the block is maintained to the extent required by the Memory Coherence Required storage attribute.

For the \(L\) operand, the value 2 is reserved. The results of executing a dcbf instruction with \(\mathrm{L}=2\) are boundedly undefined.

The function of this instruction is independent of whether the block containing the byte addressed by EA is in storage that is Write Through Required or Caching Inhibited.

This instruction is treated as a Load (see Section 4.3), except that reference and change recording need not be done.

\section*{Special Registers Altered:}

None

\section*{Extended Mnemonics:}

Extended mnemonics are provided for the Data Cache Block Flush instruction so that it can be coded with the \(L\) value as part of the mnemonic rather than as a numeric operand. These are shown as examples with the instruction. See Appendix A. "Assembler Extended Mnemonics" on page 911. The extended mnemonics are shown below.
```

Extended:
dcbf RA,RB
dcbfl RA,RB
dcbflp RA,RB

```

\section*{Equivalent to:}
```

dcbf RA,RB, 0
dcbf RA,RB, 1
dcbf RA,RB, 3

```

Except in the dcbf instruction description in this section, references to "dcbf" in Books l-III imply L=0 unless otherwise stated or obvious from context; "dcbfl" is used for \(\mathrm{L}=1\) and "dcbflp" is used for \(\mathrm{L}=3\).

\section*{Programming Note}
dcbf serves as both a basic and an extended mnemonic. The Assembler will recognize a dcbf mnemonic with three operands as the basic form, and a dcbf mnemonic with two operands as the extended form. In the extended form the \(L\) operand is omitted and assumed to be 0 .

\section*{Programming Note}
dcbf with \(L=1\) can be used to provide a hint that a block in this processor's data cache will not be reused soon.
dcbf with \(L=3\) can be used to flush a block from the processor's primary data cache but reduce the latency of a subsequent access. For example, the block may be evicted from the primary data cache but a copy retained in a lower level of the cache hierarchy.

Programs which manage coherence in software must use dcbf with \(\mathrm{L}=0\).

\subsection*{4.3.2.1 Obsolete Data Cache Instructions}

The Data Stream Touch (dst), Data Stream Touch for Store (dstst), and Data Stream Stop (dss) instructions (primary opcode 31, extended opcodes 342, 374, and 822 respectively), which were proposed for addition to the Power ISA and were implemented by some processors, must be treated as no-ops (rather than as illegal instructions).
The treatment of these instructions is independent of whether other Vector instructions are available (i.e., is independent of the contents of MSR \({ }_{\text {VEC }}\) (see Book III).

\section*{Programming Note}

These instructions merely provided hints, and thus were permitted to be treated as no-ops even on processors that implemented them.

The treatment of these instructions is independent of whether other Vector instructions are available because, on processors that implemented the instructions, the instructions were available even when other Vector instructions were not.

The extended mnemonics for these instructions were dstt, dststt, and dssall.

\subsection*{4.3.3 "or" Instruction}

\section*{"or" Cache Control Hint}
or 26,26,26
This form of or provides a hint that stores caused by preceding Store and dcbz instructions should be performed with respect to other processors and mechanisms as soon as is feasible.

\section*{Extended Mnemonics:}

Additional extended mnemonic for the or hint:

\section*{Extended:}

Equivalent to:
miso or 26,26,26
"miso" is short for "make it so."

\section*{Programming Note}

This form of the or instruction can be used to reduce latency in producer-consumer applications by requesting that modified data be made visible to other processors quickly. In this example it is assumed that the base register is GPR3.

Producer:
addi r1,r0,0x1234
sth r1,0x1000(r3) \# store data value 0x1234
lwsync \# order data store before flag store
addi r2,r0,0x0001
stb r2,0x1002(r3) \# store nonzero flag byte or r26,r26,r26 \# miso
p_10op:
lbz r2, 0x1002(r3) \# load flag byte
andi. r2,r2,0x00FF
bne p_loop \# wait for consumer to clear
\# flag

Consumer:
c_loop:
lbz r2,0x1002(r3) \# load flag byte
andi. r2, r2, 0x00FF
beq c_loop \# wait for producer to set
\# flag to nonzero
lwsync \# order flag load before \# data load
lhz r1, \(0 \times 1000(r 3)\) \# load data value
lwsync \# order data load before
\# flag store
addi r2,r0,0x0000
stb r2,0x1002(r3) \# clear flag byte
or r26,r26,r26 \# miso

\section*{Programming Note}

Warning: Other forms of or \(R x, R x, R x\) that are not described in this section and in Section 3.2 may also cause program priority to change. Use of these forms should be avoided except when software explicitly intends to alter program priority. If a no-op is needed, the preferred no-op (ori \(0,0,0\) ) should be used.

\subsection*{4.4 Copy-Paste Facility}

The Copy-Paste Facility provides an optimized block move function with the capability of moving data to or from another system. ("move" is a slight misnomer; "copy" would be more accurate, but could be confusing with respect to the copy instruction, described below.) Depending on its storage operand address, a paste may also target an accelerator, enabling a copy_first/ paste_last pair to initiate an operation on the accelerator. The authorization to access address space owned by another system or to initiate an operation on an accelerator is established though a call to the hypervisor, the details of which are outside the scope of the architecture.

Because of the inherent latency in outboard authorization and data moves, the facility is designed to amortize latency across a move group. The facility provides a plurality of copy buffers to enable pipelining and other optimizations to be applied to the group of transfers. It performs a synchronous wait for completion on the paste_last of the group. A design assumption is that interruptions and failures in the move path are very rare events, so that the unit of operation is the move group. Success or failure is reported for the move group as a whole. Accelerator invocation uses a "singleton" move group. A singleton move group is a move group consisting of a copy_first followed immediately by a paste_last.
Each transfer is specified by a copy instruction and a subsequent paste[.] instruction. A variant of the copy instruction, copy_first, identifies the beginning of the move group. copy_first activates a Move Group In Progress (MGIP) state and sets a Move Group Valid (MGV) flag. The MGV flag is reset by a cp_abort instruction, a programming error in the move group, a failure in the processing of a transfer, or the completion of the move group. The CPTOGGLE flag is set by a copy and reset by a paste[.], providing a primitive sequence checking function. A variant of the paste instruction, paste. (paste_last), signals the completion of the move group and waits for the group to complete. The result returned in CR0 indicates success or failure. paste_last deactivates MGIP and resets MGV.
Any group move that cannot be completed successfully will end with the paste_last returning "move group failed" in CRO. When a move group fails, it is the responsibility of software to break the group up into singleton move groups, each potentially repeated based on a test for success. The combination of this fail-ure-based repeating and the action of the system data storage error handler will achieve successful completion of a correctly coded sequence of transfers. For incorrectly coded sequences, breaking up the group enables error(s) to be attributed to the proper singleton move(s).

A variety of situations will cause the data storage error handler to be invoked. A transfer will fail if at least one of the two addresses does not specify main storage (vs. CSM or an accelerator) or if the copy instruction specifies an accelerator. A transfer may also fail as the result of a remote authorization error, a remote page fault, insufficient data buffering capacity, or some implemen-tation-specific data moving problem. Fatal programming errors and protocol violations are presented as exceptions to the operating system. Serviceable remote failures (e.g. page faults) are presented as exceptions to the hypervisor.
Since the buffer that holds the block until a transfer is performed is hidden state (cannot be saved and restored) and there is no way to save the state of the move group, any disruption of program execution (e.g. interrupts, event-based branch) has the potential to prevent the move group from completing successfully. The software that handles the disruption is responsible to execute cp_abort to clean up state associated with an outstanding move group if it will use the Copy/Paste Facility itself or transfer control to another program that might use the facility prior to returning control to the original program.
Transfers associated with a move group may be invalid if the associated translations are modified prior to the completion of the move group. Such conflicts will cause the move group to fail. The granularity of conflict detection is implementation-dependent and may range from simply failing a group if any tlbie[l] is received to doing a comprehensive comparison of the pending transfers to the address mapping(s) that are invalidated.

\section*{Programming Note}

A singleton move group is specified as copy_first followed immediately by paste_last with no intervening instructions. This restriction eliminates any software-related contribution to transfer failure. It is always best to avoid unnecessary instructions between the copy_first and the paste_last of a move group.

\section*{Programming Note}

A failure of a singleton move will generally be the result of a shortage of the resources required to complete the operation. When the resources are known to be shared by multiple programs, a credit-based system is frequently used to improve quality of service. If such a credit system is in use, or if the resources are not shared, the program should continually repeat the singleton move until it succeeds. However, if no credit system is in use for shared resources, it may be appropriate to apply some sort of backoff algorithm after having retried the singleton move a few times.

\section*{Programming Note}

WARNING: In rare circumstances, paste_last may falsely report successful completion when the move group is coded incorrectly. This may occur if the move group sequence includes a redundant copy_first and the sequence is interrupted just prior to the redundant copy_first. Since interrupts should be rare, any sequence that returns a false positive CRO value should fail for most executions.

\section*{Programming Note}

Once a move group encounters an error, the hardware is permitted to take shortcuts on the remainder of the group to reduce resource utilization. Software that properly breaks a failed move group into singletons and completes that emulation of the move group prior to accessing the targets of the move group can not detect such shortcuts. In general the move group itself and any interleaved instructions must constitute an idempotent collection of operations.

The instruction descriptions include an example of how the hardware can take a shortcut by not performing the transfer, e.g. if MGV=0.

Engineering Note
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{Copy} & \multicolumn{2}{|l|}{X-form} \\
\hline \multicolumn{8}{|l|}{copy RA,RB,L} \\
\hline \[
31
\] & \[
6 \quad{ }_{6} \quad \text { III }
\] & 9 \({ }^{\text {L }}\) & \[
\begin{array}{l|l}
\hline \text { L } & \text { RA } \\
0 & 11
\end{array}
\] & \({ }_{16} \mathrm{RB}\) & 21 & 774 & 1
31 \\
\hline \multicolumn{8}{|l|}{if \(\mathrm{L}=1\)} \\
\hline \multicolumn{8}{|l|}{if MGIP=0} \\
\hline \multicolumn{8}{|c|}{MGIP \(\leftarrow 1\)} \\
\hline \multicolumn{8}{|l|}{MGV \(\leftarrow 1\)} \\
\hline \multicolumn{8}{|c|}{CPTOGGLE \(¢ 0\)} \\
\hline \multicolumn{8}{|l|}{else} \\
\hline \multicolumn{8}{|l|}{MGV -0} \\
\hline \multicolumn{8}{|l|}{if (MGIP=1) \& (MGV=1) \& (CPTOGGLE=0)} \\
\hline \multicolumn{8}{|l|}{if \(\mathrm{RA}=0\) then \(\mathrm{b} \leftarrow 0\)} \\
\hline \multicolumn{8}{|l|}{else \(\quad \mathrm{b} \leftarrow\) (RA)} \\
\hline \multicolumn{8}{|l|}{\(\mathrm{EA} \leftarrow \mathrm{b}+(\mathrm{RB})\)} \\
\hline \multicolumn{8}{|l|}{copy_buffer \(\leftarrow \operatorname{mem}(E A, 128)\)} \\
\hline \multicolumn{8}{|l|}{CPTOGGLE \(ヶ 1\)} \\
\hline \multicolumn{8}{|l|}{else} \\
\hline \multicolumn{8}{|l|}{MGIP \(\leftarrow 1\)} \\
\hline \multicolumn{8}{|l|}{MGV \(¢ 0\)} \\
\hline
\end{tabular}

The copy instruction copies 128 bytes from the specified location in storage into a copy buffer, as described below. If \(L=1\), the instruction identifies the beginning of a move group.

In preparation for the copy operation, the following move group state changes are made. If \(\mathrm{L}=1\) and a move group is already in progress MGV is set to zero; otherwise, MGIP is set to 1 , MGV is set to 1 , and CPTOGGLE is set to 0 .

If \(\mathrm{MGIP}=0\), the instruction is attempting to start a sequence with a plain copy. If MGV=0, either a copy_first was issued with a move group already in progress or some programming error happened before this instruction. If CPTOGGLE=1 when neither of the previous errors occurred, this is a copy following a copy. All three of these errors will result in setting MGIP=1 and MGV=0. Otherwise, this is a good copy, and the instruction copies 128 bytes of storage into the copy buffer as follows.

Let the effective address (EA) be the sum (RAIO)+(RB).

If the storage addressed by EA is in cluster shared memory, an outboard translation/authority mechanism will be invoked in addition to the normal translation/authority checking.

The 128 bytes in storage addressed by EA is loaded into the copy buffer and CPTOGGLE is set to 1 to indicate a Copy instruction has been processed.

If the EA is not a multiple of 128, the system alignment error handler is invoked.

If the specified block is in storage that is Caching Inhibited, the system data storage error handler is invoked. If any of the sequencing errors described above have
occured and this copy is part of a singleton move, the system data storage error handler will be invoked in asociation with the execution of the related paste. instruction (if another interruption did not separate the two). If any of the sequencing errors described above have occured and this copy is part of a (non-singleton) move group, the paste. will return "move group failed" in CRO.

When successful, this instruction is treated as a Load (see Section 4.3, "Cache Management Instructions"), except that the transfer ordering is described in Section 1.7.2, "Storage Ordering of Copy/Paste-Initiated Data Transfers".
Special Registers Altered:
None

\section*{Extended Mnemonics:}

Extended mnemonics for Copy:
```

Extended:
copy RA,RB
copy_first RA,RB

```

Equivalent to:
copy RA,RB, 0 copy RA,RB, 1

\section*{- Programming Note}
copy serves as both a basic and an extended mnemonic. The Assembler will recognize a copy mnemonic with three operands as the basic form, and a copy mnemonic with two operands as the extended form. In the extended form the L operand is omitted and assumed to be 0 .
```

Paste X-form
$\begin{array}{lll}\text { paste } & R A, R B, L & (L=0 ~ R c=0) \\ \text { paste. } & R A, R B, L & (L=1 R c=1)\end{array}$

| 31 |  | /// | L | RA | RB |  | 902 | Rc |
| ---: | ---: | ---: | :--- | :--- | :--- | :--- | :--- | ---: |
| 0 |  | 6 |  | 9 | 10 | 11 | 16 | 21 |

if $(M G I P=0)|(M G V=0)|(C P T O G G L E=0)$
MGIP $\leftarrow 1$
MGV $\leftarrow 0$
else
if $R A=0$ then $b \leftarrow 0$
else $\quad \mathrm{b} \leftarrow(\mathrm{RA})$
$\mathrm{EA} \leftarrow \mathrm{b}+(\mathrm{RB})$
post-mem $(E A, 128, L) \leftarrow$ copy_buffer
CPTOGGLE $\leftarrow 0$
if $\mathrm{L}=1$
if (MGIP=1) \& (MGV=1)
wait for group completion status
if group succeeded
CR0 $\leftarrow 0 \mathrm{~b} 001\left|\mid \mathrm{XER}_{\text {S0 }}\right.$
else
$\mathrm{CR} 0 \leftarrow 0 \mathrm{~b} 000\left|\mid \mathrm{XER}_{\mathrm{SO}}\right.$
else /* this is a programming error
$\mathrm{CR} 0 \leftarrow 0 \mathrm{~b} 000\left|\mid \mathrm{XER}_{S 0}\right.$
MGIP $\leftarrow 0$
MGV $\leftarrow 0$

```

The paste instruction copies 128 bytes from the copy buffer to the specified location in storage, as described below. If \(\mathrm{L}=1\), the instruction identifies the end of a move group, and the successful completion of the instruction is contingent on the successful completion of all of the paste operations in the group.

If MGIP=0, the instruction is attempting to start a sequence with a paste. If MGV=0, some programming error happened before this instruction. If CPTOGGLE=0 when neither of the previous errors occurred, this is a paste following a paste. All three of these errors will result in setting MGIP=1 and MGV=0. Otherwise, this is a good paste, and the instruction pastes the 128 bytes from the copy buffer to storage as follows.

Let the effective address (EA) be the sum (RAIO)+(RB).

If the storage addressed by EA is in cluster shared memory, an outboard translation/authority mechanism will be invoked in addition to the normal translation/authority checking.
The contents of the copy buffer are posted to be stored in the 128 bytes of storage or sent to the accelerator addressed by EA and an indicator of whether this is a paste_last is sent along. The actual update of storage will be done asynchronously, but prior to the completion of the paste_last for the move group. If it is a
paste_last, status for all the transfers will be collected to determine the success of the move group. CPTOGGLE is set to 0 to indicate a Paste instruction has been processed.

If \(L=1\) and nothing has gone wrong with the move group so far, the processor waits for the group status to be collected, and CRO is set as follows to report the status of the move group.
\begin{tabular}{l|l}
\hline CRO & Description \\
\hline Ob000IIXER \(_{\text {SO }}\) & Move group failed. \\
\hline Ob001IIXER \(_{\text {SO }}\) & Move group successful.
\end{tabular}

If \(L \neq R c\), the instruction form is invalid.
If the EA is not a multiple of 128 , the system alignment error handler is invoked.

If the specified block is in storage that is Caching Inhibited, the system data storage error handler is invoked.

If any of the following occurs as part of a (non-singleton) move group, the paste. ending the group will return "move group failed" in CRO. If any of them occurs as part of a singleton move, the behavior will be as described for each situation.
■ For a copy / paste[.] pair, if one instruction specifies a location in CSM and the other does not specify a location in (local) main storage, the system data storage error handler will be invoked.
- For a copy / paste[.] pair, that specifies an accelerator, if the copy specifies (local) main storage and the paste[.] specifies the accelerator, the accelerator operation will be initiated (once executed as a singleton move); otherwise, the system data storage error handler will be invoked.
■ If any of the sequencing errors described above have occured, the system data storage error handler will be invoked.
- If an error from the copy description that is specified to occur when the associated paste[.] is executed occurs, the system data storage error handler is invoked.

When successful, this instruction is treated as a Store (see Section 4.3, "Cache Management Instructions"), except that the transfer ordering is described in Section 1.7.2, "Storage Ordering of Copy/Paste-Initiated Data Transfers".

\section*{Special Registers Altered:}

\section*{CRO Extended Mnemonics:}

Extended mnemonics for Paste:
\begin{tabular}{ll} 
Extended: & Equivalent to: \\
paste RA,RB & paste RA,RB,0 \\
paste_last RA,RB & paste. RA,RB,1
\end{tabular}

> Programming Note
> paste serves as both a basic and an extended mnemonic. The Assembler will recognize a paste mnemonic with three operands as the basic form, and a paste mnemonic with two operands as the extended form. In the extended form the \(L\) operand is omitted and assumed to be 0 .
\[
C P_{-} \text {Abort }
\]
X-form
cp_abort
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & \multicolumn{1}{|c|}{ I/I } & III & I/I & \multicolumn{2}{|c|}{838} & \(/\) \\
0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { MGIP } \leftarrow 0 \\
& \text { MGV } \leftarrow 0 \\
& \text { reset_CP }
\end{aligned}
\]

The \(\boldsymbol{c p}\) _abort instruction causes a group move to fail if one is in progress.
If a move group was in progress, the move group is terminated, as follows.

Any lingering state from an outstanding move group is cleaned up.
Move Group In Progress is set to false and Move Group Valid is set to false.
Special Registers Altered:
None
Engineering Note

\section*{| 4.5 Atomic Memory Operations}

The Atomic Memory Operation (AMO) facility may be used to optimize performance when many software threads are manipulating shared control structures concurrently. In such situations, accessing the shared data frequently involves transfering the data from one processor's cache to another. The latency of such transfers can become the limiting factor in the performance of some environments. Rather than moving the data to the work, AMOs move the work to the data. The mental model is of an agent consisting of an execution unit and a work queue near memory that receives atomic update requests from all the processors in the system.
Despite that AMOs are performed at memory, their function is only defined for storage that is not Caching Inhibited. This is done so that software can transparently access the same data using normal loads and stores. Since the performance advantage of AMOs derives from avoiding time of flight through cache hierarchies, software should avoid frequent mixing of normal loads and stores and AMOs to the same storage locations. AMOs are also restricted to storage that is not Guarded and storage that is not Write Through Required to limit implementation complexity.
The facility specifies a set of atomic update operations that a processor may send, accompanied by operands from GPRs, to the memory to be performed. The operations are expressed using the Load Atomic (LAT) and Store Atomic (STAT) instructions. Each of these instructions performs an atomic update operation (load followed by some manipulation and a store) on some location in storage. As a result, these instructions are considered to be both fixed-point loads and fixed-point stores, and any reference elsewhere in the architecture to fixed-point loads or fixed-point stores apply to these instructions as well, except where explicitly stated otherwise or obvious from context. For example, in order to perform an AMO, it is necessary to have both read and write access to the storage location. Another example is that the DAWR will detect a match if either Data Read or Data Write is selected. Yet another example is that a Trace interrupt will indicate both a load and a store have been executed. Barrier action will be based on whether the barrier would give a load or a store the stronger ordering. The difference between the loads and stores is simply that the loads return a result to a GPR, while the stores do not. In the RTL in the following subsections, the "lat" and "stat" functions represent the manipulations performed by the memory agent. The parameters shown are the maximum storage footprint, the maximum list of registers, and the function code that are provided to the agent. If the specified registers wrap (e.g. RT=R31 and \(R T+1=R 0\) ), the wrapping is permitted. Such an instruction is not an invalid form. Destructive encodings are also permitted (i.e. a LAT specified with RT=RA).

Except in this section, references to "atomic update" in Books I-III imply use of the Load And Reserve and Store Conditional instructions unless otherwise stated or obvious from context.

\section*{Programming Note}

The best performance for the Atomic Memory Operations will be realized when the targeted storage locations are accessed only using AMOs. If it is necessary to perform other \(\mathrm{I}=0\) loads and stores to those addresses, the result will still be correct, but performance will suffer. In such circumstances, it is not helpful to performance to flush the data to memory using dcbf.

\section*{Programming Note}

Note that the descriptions of AMO operations are Endian independent. The only effect of Endian on these operations is the obvious one that byte significance within an individual datum reflects the Endian mode.

\section*{Engineering Note}

\subsection*{4.5.1 Load Atomic}

The Atomic Loads perform an atomic update to an aligned memory location and return a value to a GPR. The manipulation performed on the memory value and the value that is returned in the GPR are determined by the function code (FC) specified by the instruction. The name of each function and its associated RTL are shown in Figure 3.

Version 3.0
\begin{tabular}{|c|c|c|c|}
\hline Function Code & GPR operands & Storage operands & Function name and RTL \\
\hline 00000 & RT, RT+1 & mem(EA,s) & ```
Fetch and Add
    t \leftarrowmem(EA, s)
    t2}\leftarrow\textrm{t}+(\textrm{RT}+1
    mem(EA,s) \leftarrow t2
    RT}\leftarrow
``` \\
\hline 00001 & RT, RT+1 & mem(EA,s) & ```
Fetch and XOR
    t}\leftarrow\operatorname{mem}(EA,s
    t2}\leftarrow\textrm{t}\oplus(\textrm{RT}+1
    mem(EA,s) \leftarrowt2
    RT}\leftarrow
``` \\
\hline 00010 & RT, RT+1 & mem(EA,s) & \[
\begin{aligned}
& \text { Fetch and } O R \\
& t \leftarrow \operatorname{mem}(E A, s) \\
& t 2 \leftarrow t \mid(R T+1) \\
& \operatorname{mem}(E A, s) \leftarrow t 2 \\
& R T \leftarrow t
\end{aligned}
\] \\
\hline 00011 & RT, RT+1 & mem(EA,s) & \[
\begin{aligned}
& \text { Fetch and AND } \\
& t \leftarrow \operatorname{mem}(E A, s) \\
& t 2 \leftarrow t \&(R T+1) \\
& \operatorname{mem}(E A, s) \leftarrow t 2 \\
& R T \leftarrow t
\end{aligned}
\] \\
\hline 00100 & RT, RT+1 & mem(EA,s) & ```
Fetch and Maximum Unsigned
    t}\leftarrow\operatorname{mem}(EA,s
    t2 \leftarrow maximum_unsigned (t, (RT+1))
    mem(EA,s) \leftarrow t2
    RT}\leftarrow
``` \\
\hline 00101 & RT, RT+1 & mem(EA,s) & ```
Fetch and Maximum Signed
    t}\leftarrow\operatorname{mem}(EA,s
    t2 \leftarrow maximum_signed(t, (RT+1))
    mem(EA,s) \leftarrow t2
    RT}\leftarrow
``` \\
\hline 00110 & RT, RT+1 & mem(EA,s) & ```
Fetch and Minimum Unsigned
    t}\leftarrow\operatorname{mem}(EA,s
    t2 \leftarrow minimum_unsigned (t, (RT+1))
    mem (EA,s) \leftarrow t2
    RT}\leftarrow
``` \\
\hline 00111 & RT, RT+1 & mem(EA,s) & ```
Fetch and Minimum Signed
    t}\leftarrow\textrm{mem}(EA,s
    t2 \leftarrow minimum_signed(t, (RT+1))
    mem(EA,s) \leftarrow t2
    RT}\leftarrow
``` \\
\hline 01000 & RT, RT+1 & mem(EA,s) & ```
Swap
    t}\leftarrow\operatorname{mem}(EA,s
    mem (EA,s) \leftarrow(RT+1)
    RT}\leftarrow
``` \\
\hline 10000 & RT, RT1, RT+2 & mem(EA,s) & ```
Compare and Swap Not Equal
    t}\leftarrow\operatorname{mem}(EA,s
    if t != (RT+1) then mem(EA,S) \leftarrow(RT+2)
    RT}\leftarrow
``` \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|}
\hline 11000 & RT & \[
\begin{aligned}
& \text { mem(EA,s) } \\
& m e m(E A+s, s)
\end{aligned}
\] & ```
Fetch and Increment Bounded
    \(t \leftarrow \operatorname{mem}(E A, s)\)
    \(\mathrm{t} 2 \leftarrow \mathrm{mem}(\mathrm{EA}+\mathrm{s}, \mathrm{s})\)
    if \(t!=t 2\) then
        \(\operatorname{mem}(E A, s) \leftarrow t+1\)
        \(R T \leftarrow t\)
    else \(R T \leftarrow 1 \ll\left(s^{*} 8-1\right)\)
``` \\
\hline 11001 & RT & \[
\begin{aligned}
& \text { mem(EA,s) } \\
& \text { mem }(E A+s, s)
\end{aligned}
\] & \[
\begin{aligned}
& \text { Fetch and Increment Equal } \\
& t \leftarrow \operatorname{mem}(E A, s) \\
& t 2 \leftarrow \operatorname{mem}(E A+s, s) \\
& \text { if } t=t 2 \text { then } \\
& \operatorname{mem}(E A, s) \leftarrow t+1 \\
& R T \leftarrow t \\
& \text { else } R T \leftarrow 1 \ll(s * 8-1)
\end{aligned}
\] \\
\hline 11100 & RT & \[
\begin{aligned}
& \text { mem(EA-s,s) } \\
& \text { mem(EA, s) }
\end{aligned}
\] & ```
Fetch and Decrement Bounded
    \(t \leftarrow \operatorname{mem}(E A, s)\)
    \(\mathrm{t} 2 \leftarrow \operatorname{mem}(E A-\mathrm{s}, \mathrm{s})\)
    if \(t\) != t2 then
        \(\operatorname{mem}(\mathrm{EA}, \mathrm{s}) \leftarrow \mathrm{t}-1\)
        \(R T \leftarrow t\)
    else \(R T \leftarrow 1 \ll\left(s^{*} 8-1\right)\)
``` \\
\hline
\end{tabular}

\section*{Notes:}
\(\mathrm{s}=\) operand size in number of bytes
Function codes not listed in this table are considered invalid.
For word atomics, only the least significant word of each source register is used, and the least significant word of the target register is updated with the result, while the upper word is set to zero.

I Figure 3. Load Atomic function codes
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{4}{|l|}{Load Word Atomic} & \multicolumn{2}{|r|}{X-form} \\
\hline Iwat & RT, & A,FC & & & \\
\hline \[
0
\] & \[
6_{6} \mathrm{RT}
\] & \[
{ }_{11} \mathrm{RA}
\] & \[
{ }_{16} \mathrm{FC}
\] & 582 & 1
31 \\
\hline
\end{tabular}
```

if RA=0 then EA }\leftarrow
else EA \leftarrow(RA)
(RT 32:63,mem (EA,4))\leftarrow lat(mem (EA-4,12), RT+1 32:63,
RT+2;22:63, FC)
RT}0:31 \leftarrow

```

Let the effective address (EA) be (RA). The least significant word of RT and the word of storage at EA are updated as specified by load atomic function code FC. The most significant word of RT is set to zero. Input operands are function code specific, and may include the least significant words of RT+1 and RT+2, and mem(EA-4,12)

Figure 3 contains the valid function codes. An attempt to execute Iwat specifying an Invalid function code will cause the system data storage error handler to be invoked.

The portion of mem(EA-4,12) accessed by the instruction must be contained within an aligned 32-byte block of storage. If it is not, the system alignment error handler will be invoked.
This instruction is treated as a Load (see Section 4.3), except with regard to storage protection, change recording, and the requirement that stores not be performed out of order.
Special Registers Altered:
None
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{4}{|l|}{Load Doubleword Atomic} & \multicolumn{2}{|r|}{\(X\)-form} \\
\hline Idat & RT, & A,FC & & & \\
\hline \[
31
\] & \[
{ }_{6} \mathrm{RT}
\] & \[
{ }_{11} \mathrm{RA}
\] & \[
{ }_{16} \mathrm{FC}
\] & \[
\begin{array}{ll}
21 & 614 \\
21
\end{array}
\] & 1
31 \\
\hline
\end{tabular}

Let the effective address (EA) be (RA). RT and the doubleword of storage at EA are updated as specified by load atomic function code FC. Input operands are function code specific, and may include RT+1, RT +2 , and mem(EA-8,24)
Figure 3 contains the valid function codes. An attempt to execute Idat specifying an Invalid function code will cause the system data storage error handler to be invoked.

The portion of mem(EA-8,24) accessed by the instruction must be contained within an aligned 32-byte block of storage. If it is not, the system alignment error handler will be invoked.

This instruction is treated as a Load (see Section 4.3), except with regard to storage protection, change recording, and the requirement that stores not be performed out of order.

\section*{Special Registers Altered:}

None

\subsection*{4.5.2 Store Atomic}

The Atomic Stores perform an atomic update to an aligned memory location. The manipulation performed on the memory value is determined by the function code (FC) specified by the instruction. The name of each function and its associated RTL are shown in Figure 4
\begin{tabular}{|c|c|c|c|}
\hline Function Code & GPR operands & Storage operands & Function name and RTL \\
\hline 00000 & RS & mem(EA,s) & \[
\begin{aligned}
& \text { Store Add } \\
& t \leftarrow \operatorname{mem}(E A, s) \\
& t 2 \leftarrow t+(R S) \\
& \operatorname{mem}(E A, s) \leftarrow t 2
\end{aligned}
\] \\
\hline 00001 & RS & mem(EA,s) & ```
Store XOR
    t}\leftarrow\operatorname{mem}(EA,s
    t2}\leftarrow\textrm{t}\oplus(\textrm{RS}
    mem(EA,s) \leftarrow t2
``` \\
\hline 00010 & RS & mem(EA,s) & ```
Store OR
    \(t \leftarrow \operatorname{mem}(E A, s)\)
    \(\mathrm{t} 2 \leftarrow \mathrm{t} \mid\) (RS)
    \(\operatorname{mem}(E A, s) \leftarrow t 2\)
``` \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|}
\hline 00011 & RS & mem(EA,s) & \[
\begin{aligned}
& \text { Store AND } \\
& t \leftarrow \operatorname{mem}(E A, s) \\
& t 2 \leftarrow t \&(R S) \\
& \operatorname{mem}(E A, s) \leftarrow t 2
\end{aligned}
\] \\
\hline 00100 & RS & mem(EA,s) & Store Maximum Unsigned
```

    t \leftarrow mem(EA, s)
    t2 \leftarrow maximum_unsigned(t, (RS))
    mem(EA,s) \leftarrow t2
    ``` \\
\hline 00101 & RS & mem(EA,s) & ```
Store Maximum Signed
    t \leftarrow mem(EA, s)
    t2 \leftarrow maximum_signed(t,(RS))
    mem(EA,s) \leftarrow t2
``` \\
\hline 00110 & RS & mem(EA,s) & ```
Store Minimum Unsigned
    t}\leftarrow\operatorname{mem}(EA,s
    t2 \leftarrow minimum_unsigned(t, (RS))
    mem(EA,s) \leftarrow t2
``` \\
\hline 00111 & RS & mem(EA,s) & ```
Store Minimum Signed
    t}\leftarrow\operatorname{mem}(EA,s
    t2 \leftarrow minimum_signed(t,(RS))
    mem (EA,s) \leftarrow t2
``` \\
\hline 11000 & RS & \[
\begin{aligned}
& \operatorname{mem}(E A, s) \\
& \operatorname{mem}(E A+s, s)
\end{aligned}
\] & \[
\begin{aligned}
& \text { Store Twin } \\
& t \leftarrow \operatorname{mem}(E A, s) \\
& t 2 \leftarrow \operatorname{mem}(E A+s, s) \\
& \text { if } t=t 2 \text { then } \\
& \quad \operatorname{mem}(E A, s) \leftarrow(\text { RS }) \\
& \operatorname{mem}(E A+s, s) \leftarrow(R S)
\end{aligned}
\] \\
\hline
\end{tabular}

\section*{Notes:}
\(\mathrm{s}=\) operand size in number of bytes
Function codes not listed in this table are considered invalid.
For word atomics, only the least significant word of each source register is used.
| Figure 4. Store Atomic function codes
Store Word Atomic
stwat
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS,RA,FC \\
\hline 0 & & RS & RA & RA & FC & \\
\hline
\end{tabular}
```

if RA=0 then EA \leftarrow0
else EA\leftarrow(RA)
mem(EA,8)\leftarrow\operatorname{stat (mem(EA,8), RS 32:63, FC)}

```

Let the effective address (EA) be (RA). Four or eight bytes of storage at EA are updated as specified by store atomic function code FC. Input operands are function code specific, and may include \(\mathrm{RS}_{32: 63}\) and mem(EA,8).
Figure 4 contains the valid function codes. An attempt to execute stwat specifying an Invalid function code will cause the system data storage error handler to be invoked.

The portion of mem(EA,8) accessed by the instruction must be contained within an aligned 32-byte block of storage. If it is not, the system alignment error handler will be invoked.

This instruction is treated as a Store (see Section 4.3).
Special Registers Altered:
None

Store Doubleword Atomic
X-form
stdat RS,RA,FC
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & RS & RA & FC & & 742 & \(/\) \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

if RA=0 then EA }\leftarrow
else EA \leftarrow (RA)
mem(EA,16)\leftarrow stat(mem(EA,16), RS, FC)

```

Let the effective address (EA) be (RA). Eight or sixteen bytes of storage at EA are updated as specified by store atomic function code FC. Input operands are function code specific, and may include RS and mem(EA,16).
Figure 4 contains the valid function codes. An attempt to execute stdat specifying an Invalid function code will cause the system data storage error handler to be invoked.

The portion of mem(EA,16) accessed by the instruction must be contained within an aligned 32-byte block of storage. If it is not, the system alignment error handler will be invoked.

This instruction is treated as a Store (see Section 4.3).
Special Registers Altered:
None

\subsection*{4.6 Synchronization Instructions}

The synchronization instructions are used to ensure that certain instructions have completed before other
instructions are initiated, or to control storage access ordering, or to support debug operations.

\subsection*{4.6.1 Instruction Synchronize Instruction}

\section*{Instruction Synchronize \\ XL-form}
isync
\begin{tabular}{|l|l|l|l|l|ll|l|}
\hline 19 & \multicolumn{1}{|c|}{\(/ / /\)} & \multicolumn{1}{c|}{\(/ / /\)} & \multicolumn{1}{c|}{\(/ / /\)} & & 150 & \(/\) \\
0 & & 6 & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}

Executing an isync instruction ensures that all instructions preceding the isync instruction have completed before the isync instruction completes, and that no subsequent instructions are initiated until after the isync instruction completes. It also ensures that all instruction cache block invalidations caused by icbi instructions preceding the isync instruction have been performed with respect to the processor executing the isync instruction, and then causes any prefetched instructions to be discarded.

Except as described in the preceding sentence, the isync instruction may complete before storage accesses associated with instructions preceding the isync instruction have been performed.
This instruction is context synchronizing (see Book III).

\section*{Special Registers Altered: \\ None}

\subsection*{4.6.2 Load and Reserve and Store Conditional Instructions}

The Load And Reserve and Store Conditional instructions can be used to construct a sequence of instructions that appears to perform an atomic update operation on an aligned storage location. See Section 1.7.4, "Atomic Update" for additional information about these instructions.

The Load And Reserve and Store Conditional instructions are fixed-point Storage Access instructions; see Section 3.3.1, "Fixed-Point Storage Access Instructions", in Book I.

The storage location specified by the Load And Reserve and Store Conditional instructions must be in storage that is Memory Coherence Required if the location may be modified by another processor or mechanism. If the specified location is in storage that is Write Through Required or Caching Inhibited, the system data storage error handler is invoked.

The Load and Reserve instructions include an Exclusive Access hint (EH), which can be used to indicate that the instruction sequence being executed is implementing one of two types of algorithms:

\section*{Atomic Update (EH=0)}

This hint indicates that the program is using a fetch and operate (e.g., fetch and add) or some similar algorithm and that all programs accessing the shared variable are likely to use a similar operation to access the shared variable for some time.

\section*{Exclusive Access (EH=1)}

This hint indicates that the program is attempting to acquire a lock and if it succeeds, will perform another store to the lock variable (releasing the lock) before another program attempts to modify the lock variable.

\section*{Programming Note}

The Memory Coherence Required attribute on other processors and mechanisms ensures that their stores to the reservation granule will cause the reservation created by the Load And Reserve instruction to be lost.

\section*{Programming Note}

Because the Load And Reserve and Store Conditional instructions have implementation dependencies (e.g., the granularity at which reservations are managed), they must be used with care. The operating system should provide system library programs that use these instructions to implement the high-level synchronization functions (Test and Set, Compare and Swap, locking, etc.; see Appendix B) that are needed by application programs. Application programs should use these library programs, rather than use the Load And Reserve and Store Conditional instructions directly.

\section*{Programming Note}
\(\mathrm{EH}=1\) should be used when the program is obtaining a lock variable which it will subsequently release before another program attempts to perform a store to it. When contention for a lock is significant, using this hint may reduce the number of times a cache block is transferred between processor caches.
EH \(=0\) should be used when all accesses to a mutex variable are performed using an instruction sequence with Load and Reserve followed by Store Conditional (e.g., emulating atomic update primitives such as "Fetch and Add;" see Appendix B). The processor may use this hint to optimize the cache to cache transfer of the block containing the mutex variable, thus reducing the latency of performing an operation such as 'Fetch and Add'.

\section*{Programming Note}

Either value of the EH field is appropriate for a Load and Reserve instruction that is intended to establish a reservation for a subsequent waitrsv and not a subsequent Store Conditional instruction.

\section*{Programming Note}

Warning: On some processors that comply with versions of the architecture that precede Version 2.00, executing a Load And Reserve instruction in which EH = 1 will cause the illegal instruction error handler to be invoked.
```

RESERVE_ADDR \leftarrow real_addr(EA)

```
\(R T \leftarrow{ }^{560}| | \operatorname{MEM}(E A, 1)\)

Let the effective address (EA) be the sum (RAIO)+(RB). The byte in storage addressed by EA is loaded into \(R T_{56: 63} . \mathrm{RT}_{0: 55}\) are set to 0 .
This instruction creates a reservation for use by a stbcx. instruction. A real address computed from the EA as described in Section 1.7.4.1 is associated with the reservation, and replaces any address previously associated with the reservation. A length of 1 byte is associated with the reservation, and replaces any length previously associated with the reservation.

The value of EH provides a hint as to whether the program will perform a subsequent store to the byte in storage addressed by EA before some other processor attempts to modify it.

0 Other programs might attempt to modify the byte in storage addressed by EA regardless of the result of the corresponding stbcx. instruction.
1 Other programs will not attempt to modify the byte in storage addressed by EA until the program that has acquired the lock performs a subsequent store releasing the lock.

\section*{Special Registers Altered:}

None

\section*{Programming Note}

Ibarx serves as both a basic and an extended mnemonic. The Assembler will recognize a Ibarx mnemonic with four operands as the basic form, and a lbarx mnemonic with three operands as the extended form. In the extended form the EH operand is omitted and assumed to be 0 .

\section*{Load Byte And Reserve Indexed X-form}

Ibarx RT,RA,RB,EH
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & & RT & RA & RB & \\
\hline 0 & & 6 & & 11 & 16 \\
\hline
\end{tabular}
```

if RA = 0 then b \& 0
else }\quad\textrm{b}\leftarrow(RA
EA}\leftarrow\textrm{b}+(\textrm{RB}
RESERVE }\leftarrow
RESERVE_LENGTH }\leftarrow

```

\section*{Load Halfword And Reserve Indexed X-form}
\begin{tabular}{|c|c|c|c|c|c|}
\hline \multicolumn{6}{|l|}{Inarx RT,RA,RB,EH} \\
\hline \[
0_{0}^{31}
\] & \({ }_{6}\) RT & \[
{ }_{11} \text { RA }
\] & \[
{ }_{16} \mathrm{RB}
\] & \[
21 \quad 116
\] & \begin{tabular}{|c|}
EH \\
31
\end{tabular} \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(RA
EA \leftarrow b + (RB)
RESERVE }\leftarrow
RESERVE_LENGTH }\leftarrow
RESERVE_ADDR \leftarrow real_addr(EA)
RT \leftarrow 480 || MEM(EA, 2)

```

Let the effective address (EA) be the sum (RAIO)+(RB). The halfword in storage addressed by EA is loaded into \(\mathrm{RT}_{48: 63} . \mathrm{RT}_{0: 47}\) are set to 0 .

This instruction creates a reservation for use by a sthcx. instruction. A real address computed from the EA as described in Section 1.7.4.1 is associated with the reservation, and replaces any address previously associated with the reservation. A length of 2 bytes is associated with the reservation, and replaces any length previously associated with the reservation.

The value of EH provides a hint as to whether the program will perform a subsequent store to the halfword in storage addressed by EA before some other processor attempts to modify it.

0 Other programs might attempt to modify the halfword in storage addressed by EA regardless of the result of the corresponding sthcx. instruction.
1 Other programs will not attempt to modify the halfword in storage addressed by EA until the program that has acquired the lock performs a subsequent store releasing the lock.

EA must be a multiple of 2 . If it is not, either the system alignment error handler is invoked or the results are boundedly undefined.

\section*{Special Registers Altered:}

None

\section*{Programming Note}

Iharx serves as both a basic and an extended mnemonic. The Assembler will recognize a Iharx mnemonic with four operands as the basic form, and a Iharx mnemonic with three operands as the extended form. In the extended form the EH operand is omitted and assumed to be 0 .

Load Word And Reserve Indexed X-form
Iwarx RT,RA,RB,EH
\begin{tabular}{|l|l|l|c|c|c|c|}
\hline 31 & RT & RA & RB & \multicolumn{2}{|c|}{20} & EH \\
0 & & & 6 & 11 & 16 & 21 \\
\hline
\end{tabular}
```

if $R A=0$ then $b \leftarrow 0$
else $\quad \mathrm{b} \leftarrow(\mathrm{RA})$
$\mathrm{EA} \leftarrow \mathrm{b}+(\mathrm{RB})$
RESERVE $\leftarrow 1$
RESERVE_LENGTH $\leftarrow 4$
RESERVE_ADDR $\leftarrow$ real_addr $(E A)$
$\mathrm{RT} \leftarrow 3^{-} 0| | \operatorname{MEM}(E A, 4)$

```

Let the effective address (EA) be the sum (RAl0)+(RB). The word in storage addressed by EA is loaded into \(\mathrm{RT}_{32: 63} . \mathrm{RT}_{0: 31}\) are set to 0 .

This instruction creates a reservation for use by a stwcx. instruction. A real address computed from the EA as described in Section 1.7.4.1 is associated with the reservation, and replaces any address previously associated with the reservation. A length of 4 bytes is associated with the reservation, and replaces any length previously associated with the reservation.

The value of EH provides a hint as to whether the program will perform a subsequent store to the word in storage addressed by EA before some other processor attempts to modify it.

0 Other programs might attempt to modify the word in storage addressed by EA regardless of the result of the corresponding stwcx. instruction.
1 Other programs will not attempt to modify the word in storage addressed by EA until the program that has acquired the lock performs a subsequent store releasing the lock.

EA must be a multiple of 4 . If it is not, either the system alignment error handler is invoked or the results are boundedly undefined.

\section*{Special Registers Altered: None}

\section*{Programming Note}

Iwarx serves as both a basic and an extended mnemonic. The Assembler will recognize a Iwarx mnemonic with four operands as the basic form, and a Iwarx mnemonic with three operands as the extended form. In the extended form the EH operand is omitted and assumed to be 0 .

\section*{Store Byte Conditional Indexed X-form}
stbcx. RS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{2}{|c|}{ RS } & RA & RB & & 694 \\
\hline 0 & & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

if RA = 0 then b \leftarrow
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
if RESERVE then
if RESERVE_LENGTH = 1 then
if RESERVE_ADDR = real_addr(EA) then
MEM(EA, 1) \& (RS) 56:63
undefined_case }\leftarrow
store_performed \leftarrow \&
else
z }\leftarrow\mathrm{ smallest real page size supported by
implementation
if RESERVE_ADDR }\div\textrm{z}= real_addr (EA) \div z then
undefined_case < <
else
undefined_case \leftarrow0
store_performed \leftarrow0
else
undefined_case \leftarrow 1
else
undefined_case }\leftarrow
store_performed \leftarrow0
if undefined_case then
u1 }\leftarrow\mathrm{ undefined 1-bit value
if u1 then
MEM(EA, 1)}\leftarrow(\textrm{RS}\mp@subsup{)}{56:63}{
u2 }\leftarrow\mathrm{ undefined 1-bit value
CRO \leftarrow0.b00 || u2 || XER SO
else
CR0 \leftarrow 0b00 || store_performed || XER SO
RESERVE }\leftarrow

```

Let the effective address (EA) be the sum (RAIO)+(RB).
If a reservation exists, the length associated with the reservation is 1 byte, and the real storage location specified by the stbcx. is the same as the real storage location specified by the Ibarx instruction that established the reservation, (RS) \(56: 63\) are stored into the byte in storage addressed by EA.
If a reservation exists, the length associated with the reservation is 1 byte, and the real storage location specified by the stbcx. is not the same as the real storage location specified by the Ibarx instruction that established the reservation, the following applies.
■ Let z denote an aligned block of real storage whose size is the smallest real page size supported by the implementation. If the real storage location specified by the stbcx. is in the same \(z\) as the real storage location specified by the Ibarx instruction that established the reservation, it is undefined whether \((\mathrm{RS})_{56: 63}\) are stored into the byte in storage addressed by EA. Otherwise, no store is performed.

If a reservation exists and the length associated with the reservation is not 1 byte, it is undefined whether \((\mathrm{RS})_{56: 63}\) are stored into the byte in storage addressed by EA.

If a reservation does not exist, no store is performed.
CR Field 0 is set as follows. n is a 1-bit value that indicates whether the store was performed, except that if, per the preceding description, it is undefined whether the store is performed, the value of n is undefined (and need not reflect whether the store was performed).
\(\mathrm{CRO}_{\text {LT GT EQ SO }}=0 \mathrm{bOO}\|\mathrm{n}\|\) XER \(_{\text {SO }}\)
The reservation is cleared.
Special Registers Altered:
CRO

\section*{Store Halfword Conditional Indexed X-form}
sthcx. RS,RA,RB

RESERVE \(\leftarrow 0\)

Let the effective address (EA) be the sum (RAIO)+(RB).
If a reservation exists, the length associated with the reservation is 2 bytes, and the real storage location specified by the sthcx. is the same as the real storage location specified by the Iharx instruction that established the reservation, \((\mathrm{RS})_{48: 63}\) are stored into the halfword in storage addressed by EA.

If a reservation exists, the length associated with the reservation is 2 bytes, and the real storage location specified by the sthcx. is not the same as the real storage location specified by the Iharx instruction that established the reservation, the following applies.
■ Let z denote an aligned block of real storage whose size is the smallest real page size supported by the implementation. If the real storage location specified by the sthcx. is in the same \(z\) as the real storage location specified by the Iharx instruction that established the reservation, it is undefined whether \((\mathrm{RS})_{48: 63}\) are stored into the halfword in storage addressed by EA. Otherwise, no store is performed.

If a reservation exists and the length associated with the reservation is not 2 bytes, it is undefined whether \((\mathrm{RS})_{48: 63}\) are stored into the halfword in storage addressed by EA.

If a reservation does not exist, no store is performed.
CR Field 0 is set as follows. n is a 1-bit value that indicates whether the store was performed, except that if, per the preceding description, it is undefined whether the store is performed, the value of \(n\) is undefined (and need not reflect whether the store was performed).
\[
\mathrm{CRO}_{\text {LT GT EQ }} \mathrm{SO}=0 \mathrm{~b} 00\|\mathrm{l}\| \mathrm{XER}_{\text {SO }}
\]

The reservation is cleared.
EA must be a multiple of 2 . If it is not, either the system alignment error handler is invoked or the results are boundedly undefined.

\section*{Special Registers Altered:}

CRO

\section*{Store Word Conditional Indexed X-form}
stwcx. RS,RA,RB
\begin{tabular}{|l|l|l|l|c|c|c|}
\hline 31 & \multicolumn{1}{|c|}{ RS } & RA & RB & & 150 & 1 \\
0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

if RA = 0 then b \leftarrow
else b}\leftarrow(RA
EA \leftarrow b + (RB)
if RESERVE then
if RESERVE_LENGTH = 4 then
if RESERVE_ADDR = real_addr(EA) then
MEM (EA, 4) \leftarrow(RS) 32:63
undefined_case }\leftarrow
store_performed \leftarrow <
else
z \leftarrow smallest real page size supported by
implementation
if RESERVE_ADDR }\div\textrm{z}= real_addr (EA) \div z then
undefined_case \leftarrow 1
else
undefined_case }\leftarrow
store_performed \leftarrow0
else
undefined_case \leftarrow 1
else
undefined_case }\leftarrow
store_performed \leftarrow0
if undefined_case then
u1 }\leftarrow\mathrm{ undefined 1-bit value
if u1 then
MEM(EA,4)}\leftarrow(\textrm{RS}\mp@subsup{)}{32:63}{
u2 }\leftarrow\mathrm{ undefined 1-bit value
CRO \leftarrow0.b00 || u2 || XER SO
else
CR0 \leftarrow 0b00 || store_performed || XER SO
RESERVE }\leftarrow

```

Let the effective address (EA) be the sum (RAIO)+(RB).
If a reservation exists, the length associated with the reservation is 4 bytes, and the real storage location specified by the stwcx. is the same as the real storage location specified by the Iwarx instruction that established the reservation, \((R S)_{32: 63}\) are stored into the word in storage addressed by EA.
If a reservation exists, the length associated with the reservation is 4 bytes, and the real storage location specified by the stwcx. is not the same as the real storage location specified by the Iwarx instruction that established the reservation, the following applies.
- Let \(z\) denote an aligned block of real storage whose size is the smallest real page size supported by the implementation. If the real storage location specified by the stwcx. is in the same \(z\) as the real storage location specified by the Iwarx instruction that established the reservation, it is undefined whether \((\mathrm{RS})_{32: 63}\) are stored into the word in storage addressed by EA. Otherwise, no store is performed.

If a reservation exists and the length associated with the reservation is not 4 bytes, it is undefined whether \((\mathrm{RS})_{32: 63}\) are stored into the word in storage addressed by EA.

If a reservation does not exist, no store is performed.
CR Field 0 is set as follows. n is a 1-bit value that indicates whether the store was performed, except that if, per the preceding description, it is undefined whether the store is performed, the value of \(n\) is undefined (and need not reflect whether the store was performed).
\[
\mathrm{CRO}_{\text {LT GT EQ SO }}=0 \mathrm{~b} 00\|\mathrm{n}\| \mathrm{XER}_{\text {SO }}
\]

The reservation is cleared.
EA must be a multiple of 4 . If it is not, either the system alignment error handler is invoked or the results are boundedly undefined.

\section*{Special Registers Altered:}

CRO

\subsection*{4.6.2.1 64-Bit Load and Reserve and Store Conditional Instructions}

\section*{Load Doubleword And Reserve Indexed X-form}

Idarx RT,RA,RB,EH
\begin{tabular}{|l|l|l|l|c|c|c|r|}
\hline 31 & RT & RA & RB & & 84 & EH \\
0 & & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else }\quad\textrm{b}\leftarrow(\textrm{RA}
EA}\leftarrow\textrm{b}+(\textrm{RB}
RESERVE \leftarrow1
RESERVE_LENGTH \leftarrow 8
RESERVE_ADDR \leftarrow real_addr(EA)
RT \leftarrow MEM(EA, 8)

```

Let the effective address (EA) be the sum (RAIO)+(RB). The doubleword in storage addressed by EA is loaded into RT.

This instruction creates a reservation for use by a stdcx. instruction. A real address computed from the EA as described in Section 1.7.4.1 is associated with the reservation, and replaces any address previously associated with the reservation. A length of 8 bytes is associated with the reservation, and replaces any length previously associated with the reservation.

The value of EH provides a hint as to whether the program will perform a subsequent store to the doubleword in storage addressed by EA before some other processor attempts to modify it.

0 Other programs might attempt to modify the doubleword in storage addressed by EA regardless of the result of the corresponding stdcx. instruction.
1 Other programs will not attempt to modify the doubleword in storage addressed by EA until the program that has acquired the lock performs a subsequent store releasing the lock.

EA must be a multiple of 8 . If it is not, either the system alignment error handler is invoked or the results are boundedly undefined.
Special Registers Altered:
None

\section*{Programming Note}

Idarx serves as both a basic and an extended mnemonic. The Assembler will recognize a Idarx mnemonic with four operands as the basic form, and a Idarx mnemonic with three operands as the extended form. In the extended form the EH operand is omitted and assumed to be 0 .

\section*{Store Doubleword Conditional Indexed X-form}
stdcx. RS,RA,RB
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & RS & RA & RB & & 214 & 1 \\
0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

if $\mathrm{RA}=0$ then $\mathrm{b} \leftarrow 0$
else $\quad \mathrm{b} \leftarrow(\mathrm{RA})$
$\mathrm{EA} \leftarrow \mathrm{b}+(\mathrm{RB})$
if RESERVE then
if RESERVE_LENGTH = 8 then
if RESERVE_ADDR = real_addr(EA) then
$\operatorname{MEM}(E A, 8) \leftarrow(\operatorname{RS})$
undefined_case $\leftarrow 0$
store_performed $\leftarrow 1$
else
$z \leftarrow$ smallest real page size supported by
implementation
if RESERVE_ADDR $\div z=$ real_addr $(E A) \div z$ then
undefined_case $\leftarrow 1$
else
undefined_case $\leftarrow 0$
store_performed $\leftarrow 0$
else
undefined_case $\leftarrow 1$
else
undefined_case $\leftarrow 0$
store_performed $\leftarrow 0$
if undefined_case then
u1 $\leftarrow$ undefined 1-bit value
if u1 then
$\operatorname{MEM}(E A, 8) \leftarrow(R S)$
u2 $\leftarrow$ undefined 1-bit value
$\mathrm{CRO} \leftarrow 0 \mathrm{~b} 00||\mathrm{u} 2|| \mathrm{XER}_{\mathrm{SO}}$
else
$\mathrm{CRO} \leftarrow 0 \mathrm{~b} 00\left|\mid\right.$ store_performed || $\mathrm{XER}_{\text {SO }}$
RESERVE $\leftarrow 0$

```

Let the effective address (EA) be the sum (RAIO)+(RB).
If a reservation exists, the length associated with the reservation is 8 bytes, and the real storage location specified by the stdcx. is the same as the real storage location specified by the Idarx instruction that established the reservation, (RS) is stored into the doubleword in storage addressed by EA.
If a reservation exists, the length associated with the reservation is 8 bytes, and the real storage location specified by the stdcx. is not the same as the real storage location specified by the Idarx instruction that established the reservation, the following applies.

■ Let z denote an aligned block of real storage whose size is the smallest real page size supported by the implementation. If the real storage location specified by the stdcx. is in the same \(z\) as the real storage location specified by the Idarx instruction that established the reservation, it is

\section*{Version 3.0}
undefined whether (RS) is stored into the doubleword in storage addressed by EA. Otherwise, no store is performed.

If a reservation exists and the length associated with the reservation is not 8 bytes, it is undefined whether (RS) is stored into the doubleword in storage addressed by EA.

If a reservation does not exist, no store is performed.
CR Field 0 is set as follows. \(n\) is a 1-bit value that indicates whether the store was performed, except that if, per the preceding description, it is undefined whether the store is performed, the value of n is undefined (and need not reflect whether the store was performed).
\[
\mathrm{CRO}_{\mathrm{LT}} \mathrm{GT} \text { EQ SO }=0 \mathrm{bOO}\|\mathrm{n}\| \mathrm{XER}_{\mathrm{SO}}
\]

The reservation is cleared.
EA must be a multiple of 8 . If it is not, either the system alignment error handler is invoked or the results are boundedly undefined.

\section*{Special Registers Altered:}

CRO

\subsection*{4.6.2.2 128-bit Load and Reserve Store Conditional Instructions}

For Iqarx, the quadword in storage addressed by EA is loaded into an even-odd pair of GPRs as follows. In Big-Endian mode, the even-numbered GPR is loaded with the doubleword from storage addressed by EA and the odd-numbered GPR is loaded with the doubleword addressed by EA+8. In Little-Endian mode, the even-numbered GPR is loaded with the byte-reversed doubleword from storage addressed by EA+8 and the odd-numbered GPR is loaded with the byte-reversed doubleword addressed by EA.

In the preferred form of the Load Quadword instruction \(R A \neq R T p+1\) and \(R B \neq R T p+1\).

For stqcx., the contents of an even-odd pair of GPRs is stored into the quadword in storage addressed by EA as follows. In Big-Endian mode, the even-numbered GPR is stored into the doubleword in storage addressed by EA and the odd-numbered GPR is stored into the doubleword addressed by EA+8. In Lit-tle-Endian mode, the even-numbered GPR is stored byte-reversed into the doubleword in storage addressed by EA+8 and the odd-numbered GPR is stored byte-reversed into the doubleword addressed by EA.

\section*{Load Quadword And Reserve Indexed X-form}
```

Iqarx RTp,RA,RB,EH
if RA = 0 then b b <0
else }\quad\textrm{b}\leftarrow(\mathrm{ RA )
EA}\leftarrow\textrm{b}+(\textrm{RB}
RESERVE \leftarrow 1
RESERVE_LENGTH \leftarrow 16
RESERVE_ADDR \leftarrow real_addr(EA)
RTp \leftarrow MEM (EA, 16)

```
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & RTp & RA & RB & 276 & EH \\
\hline 0 & & & & & 31 \\
\hline
\end{tabular}

Let the effective address (EA) be the sum (RAIO)+(RB). The quadword in storage addressed by EA is loaded into RTp.

This instruction creates a reservation for use by a stqcx. instruction. A real address computed from the EA as described in Section 1.7.4.1 is associated with the reservation, and replaces any address previously associated with the reservation. A length of 16 bytes is associated with the reservation, and replaces any length previously associated with the reservation.
The value of EH provides a hint as to whether the program will perform a subsequent store to the doubleword in storage addressed by EA before some other processor attempts to modify it.

0 Other programs might attempt to modify the doubleword in storage addressed by EA regardless of the result of the corresponding stqcx. instruction.
1 Other programs will not attempt to modify the doubleword in storage addressed by EA until the program that has acquired the lock performs a subsequent store releasing the lock.

EA must be a multiple of 16 . If it is not, either the system alignment error handler is invoked or the results are boundedly undefined.

If RTp is odd, \(\mathrm{RTp}=\mathrm{RA}\), or \(\mathrm{RTp}=\mathrm{RB}\) the instruction form is invalid. If \(R T p=R A\) or \(R T p=R B\), an attempt to execute this instruction will invoke the system illegal instruction error handler. (The RTp=RA case includes the case of \(R T p=R A=0\).)

\section*{Special Registers Altered:}

None

\section*{Programming Note}

Iqarx serves as both a basic and an extended mnemonic. The Assembler will recognize a Iqarx mnemonic with four operands as the basic form, and a Iqarx mnemonic with three operands as the extended form. In the extended form the EH operand is omitted and assumed to be 0 .

\section*{Store Quadword Conditional Indexed X-form}
stqcx. RSp,RA,RB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & RSp & RA & RB & & 182 \\
0 & & 6 & 11 & 16 & 21 \\
\hline
\end{tabular}
```

if RA = 0 then b \leftarrow0
else b
EA \leftarrow b + (RB)
if RESERVE then
if RESERVE_LENGTH = 16 then
if RESERVE_ADDR = real_addr(EA) then
MEM (EA, 16) \leftarrow (RSp)
undefined_case \leftarrow 0
store_performed \leftarrow <
else
z}\leftarrow\mathrm{ smallest real page size supported by
implementation
if RESERVE_ADDR \div z = real_addr (EA) \div z then
undefined_case }\leftarrow
else
undefined_case }\leftarrow
store_performed \leftarrow0
else
undefined_case \leftarrow 1
else
undefined_case }\leftarrow
store_performed \leftarrow 0
if undefined_case then
u1 }\leftarrow\mathrm{ undefined 1-bit value
if ul then
MEM(EA, 16) \leftarrow(RSp)
u2 \leftarrow undefined 1-bit value
CR0}\leftarrow0.\textrm{b}00||\textrm{u}2||\mp@subsup{\textrm{XER}}{\mathrm{ SO}}{
else
CRO \leftarrow 0.b00 || store_performed || XER SO
RESERVE \leftarrow 0

```

Let the effective address (EA) be the sum (RAIO)+(RB).
If a reservation exists, the length associated with the reservation is 16 bytes, and the real storage location specified by the stqcx. is the same as the real storage location specified by the Iqarx instruction that established the reservation, (RSp) is stored into the quadword in storage addressed by EA.
If a reservation exists, the length associated with the reservation is 16 bytes, and the real storage location specified by the stqcx. is not the same as the real storage location specified by the Iqarx instruction that established the reservation, the following applies.
- Let z denote an aligned block of real storage whose size is the smallest real page size supported by the implementation. If the real storage location specified by the stqcx. is in the same \(z\) as the real storage location specified by the Iqarx instruction that established the reservation, it is undefined whether ( \(R S p\) ) is stored into the quadword in storage addressed by EA. Otherwise, no store is performed.

If a reservation exists and the length associated with the reservation is not 16 bytes, it is undefined whether (RSp) is stored into the quadword in storage addressed by EA.

If a reservation does not exist, no store is performed.
CR Field 0 is set as follows. n is a 1 -bit value that indicates whether the store was performed, except that if, per the preceding description, it is undefined whether the store is performed, the value of n is undefined (and need not reflect whether the store was performed).
\[
\mathrm{CRO}_{\mathrm{LT}} \text { GT EQ SO }=0 \mathrm{~b} 00\|\mathrm{n}\| \mathrm{XER}_{\text {SO }}
\]

The reservation is cleared.
EA must be a multiple of 16 . If it is not, either the system alignment error handler is invoked or the results are boundedly undefined.
If RSp is odd, the instruction form is invalid.

\section*{Special Registers Altered:}

CRO

\subsection*{4.6.3 Memory Barrier Instructions}

The Memory Barrier instructions can be used to control the order in which storage accesses and move groups of data transfers are performed. See Section 1.8, "Transactions" for a description of how the Memory Bar-
rier instructions interact with transactions. Additional information about these instructions and about related aspects of storage management can be found in Book III.

\section*{Synchronize}
\(X\)-form
```

sync L

```
\begin{tabular}{|l|l|l|l|l|l|l|l|l|}
\hline 31 & I/I & L & & I & & I/I & & 598 \\
\hline 0 & & 6 & 9 & 11 & & 16 & & 21 \\
& & & \\
\hline
\end{tabular}
```

switch(L)
case(0): hwsync
case(1): lwsync
case(2): ptesync

```

The sync instruction creates a memory barrier (see Section 1.7.1). The set of storage accesses and/or move groups of data transfers that is ordered by the memory barrier depends on the contents of the \(L\) field as follows.
- L=0 ("heavyweight sync")

The memory barrier provides an ordering function for the storage accesses and move groups of data transfers associated wth all instructions that are executed by the processor executing the sync instruction. The applicable pairs are all pairs \(\mathrm{a}_{\mathrm{i}}, \mathrm{b}_{\mathrm{j}}\) of storage accesses and move groups in which \(b_{j}\) is a data access or move group, except that if \(a_{i}\) is the storage access caused by an icbi instruction then \(b_{j}\) may be performed with respect to the processor executing the sync instruction before \(a_{i}\) is performed with respect to that processor.

■ L=1 ("lightweight sync")
The memory barrier provides an ordering function for the storage accesses caused by Load, Store, and dcbz instructions that are executed by the processor executing the sync instruction and for which the specified storage location is in storage that is Memory Coherence Required and is neither Write Through Required nor Caching Inhibited. The applicable pairs are all pairs \(\mathrm{a}_{\mathrm{i}}, \mathrm{b}_{\mathrm{j}}\) of storage accesses except those in which \(\mathrm{a}_{\mathrm{i}}\) is an access caused by a Store or \(d c b z\) instruction and \(b_{j}\) is an access caused by a Load instruction.

■ L=2 ("ptesync")
The set of storage accesses that is ordered by the memory barrier is described in

Section 5.9.2 of Book III, as are additional properties of the sync instruction with \(\mathrm{L}=2\).

The ordering done by the memory barrier is cumulative (regardless of \(L\) value).
If \(L=0\) (or \(L=2\) ), the sync instruction has the following additional properties.
■ Executing the sync instruction ensures that all instructions preceding the sync instruction have completed before the sync instruction completes, and that no subsequent instructions are initiated until after the sync instruction completes.
■ The sync instruction is execution synchronizing (see Book III). However, address translation and reference and change recording (see Book III) associated with subsequent instructions may be performed before the sync instruction completes.
■ The memory barrier provides the additional ordering function such that if a given instruction that is the result of a store in set B is executed, all applicable storage accesses in set A have been performed with respect to the processor executing the instruction to the extent required by the associated memory coherence properties. The single exception is that any storage access in set A that is caused by an icbi instruction executed by the processor executing the sync instruction (P1) may not have been performed with respect to P1 (see the description of the icbi instruction on page 842).
The cumulative properties of the memory barrier apply to the execution of the given instruction as they would to a load that returned a value that was the result of a store in set B.

\section*{Programming Note}

Section 1.10 contains a detailed description of how to modify instructions such that a well-defined result is obtained.

The value \(\mathrm{L}=3\) is reserved.
The sync instruction may complete before storage accesses associated with instructions preceding the sync instruction have been performed.

\section*{Special Registers Altered:}

None

\section*{Version 3.0}

\section*{Extended Mnemonics:}

Extended mnemonics for Synchronize:
\begin{tabular}{ll} 
Extended: & Equivalent to: \\
sync & sync 0 \\
lwsync & sync 1 \\
ptesync & sync 2
\end{tabular}

Except in the sync instruction description in this section, references to "sync" in Books l-III imply \(\mathrm{L}=0\) unless otherwise stated or obvious from context; the appropriate extended mnemonics are used when other L values are intended.

\section*{Programming Note}
sync serves as both a basic and an extended mnemonic. Assemblers will recognize a sync mnemonic with one operand as the basic form, and a sync mnemonic with no operand as the extended form. In the extended form the \(L\) operand is omitted and assumed to be 0 .

\section*{Programming Note}

The sync instruction can be used to ensure that all stores into a data structure, caused by Store instructions executed in a "critical section" of a program, will be performed with respect to another processor before the store that releases the lock is performed with respect to that processor; see Section B.2, "Lock Acquisition and Release, and Related Techniques" on page 915.

The memory barrier created by a sync instruction with \(L=1\) does not order implicit storage accesses or instruction fetches. The memory barrier created by a sync instruction with \(\mathrm{L}=0\) (or \(\mathrm{L}=2\) ) orders implicit storage accesses and instruction fetches associated with instructions preceding the sync instruction but not those associated with instructions following the sync instruction.

In order to obtain the best performance across the widest range of implementations, the programmer should use the sync instruction with \(L=1\), or the eieio instruction, if any of these is sufficient for his needs; otherwise he should use sync with \(\mathrm{L}=0\). sync with \(\mathrm{L}=2\) should not be used by application programs.

\section*{Programming Note}

The functions provided by sync with \(L=1\) are a strict subset of those provided by sync with \(\mathrm{L}=0\). (The functions provided by sync with \(L=2\) are a strict superset of those provided by sync with \(L=0\); see Book III.)

Enforce In-order Execution of I/OX-form
eieio
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 31 & I/I & I/I & I/I & & 854 & 1 & \\
\hline 0 & 6 & 11 & 16 & 21 & & & 31 \\
\hline
\end{tabular}

The eieio instruction creates a memory barrier (see Section 1.7.1, "Storage Access Ordering"), which provides an ordering function for the storage accesses caused by Load, Store, and dcbz instructions executed by the processor executing the eieio instruction. These storage accesses are divided into the two sets listed below. The storage access caused by a dcbz instruction is ordered as a store.
1. Loads and stores to storage that is both Caching Inhibited and Guarded, and stores to main storage caused by stores to storage that is Write Through Required.
The applicable pairs are all pairs \(\mathrm{a}_{\mathrm{i}}, \mathrm{b}_{\mathrm{j}}\) of such accesses.
2. Stores to storage that is Memory Coherence Required and is neither Write Through Required nor Caching Inhibited.

The applicable pairs are all pairs \(\mathrm{a}_{\mathrm{i}}, \mathrm{b}_{\mathrm{j}}\) of such accesses.

The operations caused by the stream variants of the dcbt and dcbtst instructions (i.e., the providing of
I hints) are ordered by eieio as a third set of operations, the operations caused by tlbie and tlbsync instructions (see Book III) are ordered by eieio as a fourth set of operations, and the operations caused by slbieg and slbsync instructions (see Book III) are ordered by eieio as a fifth set of operations.

Each of the five sets of storage accesses or operations is ordered independently of the other four sets. The ordering done by eieio's memory barrier for the second set is cumulative; the ordering done by eieio's memory I barrier for the other four sets is not cumulative.

The eieio instruction may complete before storage accesses associated with instructions preceding the eieio instruction have been performed. The eieio instruction may complete before operations caused by \(\boldsymbol{d c b t}\) and dcbtst instructions preceding the eieio instruction have been performed

\section*{Special Registers Altered: \\ None}

\section*{Programming Note}

The eieio instruction is intended for use in doing memory-mapped I/O). Because loads, and separately stores, to storage that is both Caching Inhibited and Guarded are performed in program order (see Section 1.7.1, "Storage Access Ordering" on page 818), eieio is needed for such storage only when loads must be ordered with respect to stores.

For the eieio instruction, accesses in set \(1, a_{i}\) and \(b_{j}\) need not be the same kind of access or be to storage having the same storage control attributes. For example, \(a_{i}\) can be a load to Caching Inhibited, Guarded storage, and \(b_{j}\) a store to Write Through Required storage.

If stronger ordering is desired than that provided by eieio, the sync instruction must be used, with the appropriate value in the \(L\) field.

\section*{Programming Note}

The functions provided by eieio for its second set are a strict subset of those provided by sync with \(\mathrm{L}=1\).

\subsection*{4.6.4 Wait Instruction}

The wait instruction is used to stop instruction fetching and execution until certain events occur. These events
include exceptions (see Section 1.2.1 of Book III) and event-based branch exceptions (see Section 1.1).

\section*{Wait}

X-form
wait WC
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & I/I & WC & I/I & \multicolumn{1}{|l|}{} \\
\hline 0 & & 6 & 9 & 11 & 16 & 21 \\
\hline
\end{tabular}

The wait instruction causes instruction fetching and execution to be suspended. Instruction fetching and execution are resumed when the events specified by the WC field occur.

The values of the WC field are as follows.

\section*{Ob00}

Resume instruction fetching and execution when an exception or an event-based branch exception occurs.

0b01:11 Reserved.
The exception or EBB exception causes the wait instruction to complete and instruction fetching to resume.

When the wait instruction completes, processing is resumed either at the instruction following the wait (if interrupts and/or event-based branches are disabled) or in the corresponding interrupt or event-based branch handler (if interrupts and/or event-based branches are enabled). If an interrupt or event-based branch causes resumption of instruction execution, the interrupt or event-based branch handler will return to the instruction after the wait.

\section*{Special Registers Altered:}

None

\section*{Extended Mnemonics:}

Examples of extended mnemonics for wait.
\begin{tabular}{ll} 
Extended: & \multicolumn{2}{l}{ Equivalent to: } \\
wait & wait 0
\end{tabular}

\section*{Programming Note}

Applications that execute wait in order to suspend processing until an external event-based branch exception occurs (see Section 7.2) should enable external event-based branch exceptions (by setting \(\mathrm{BESCR}_{E E}=1\) ) and disable event-based branches (by setting \(\mathrm{BESCR}_{\mathrm{GE}}=0\) ) before executing wait. If \(\mathrm{BESCR}_{\mathrm{GE}}=1\), then the expected event-based branch exception may cause the corresponding event-based branch to occur immediately prior to execution of the wait instruction. This will result in a hang condition since the EBB exception that was expected to cause wait to resume will have already occurred.

Since exceptions corresponding to system-caused interrupts (see Section 6.4 of Book III) may occur at any time, including immediately prior to the wait instruction, applications should not depend on them to cause wait to resume. In order to ensure timely resumption, therefore, applications should execute wait only in order to suspend processing until an event-based branch exception occurs.

Also, since exceptions corresponding to interrupts can cause wait to resume at any time without any EBB exception having occurred, programs that execute wait should check that the expected condition has actually occurred after the wait instruction completes. If the expected condition has not occurred, wait should be re-executed. An example code usage is shown below.
while ( \(\urcorner\) expected condition), wait

\section*{I}

\section*{Programming Note}

The wait instruction frees computational resources which might be allocated to another program or converted into power savings.

\title{
Chapter 5. Transactional Memory Facility
}

\subsection*{5.1 Transactional Memory Facility Overview}

This chapter describes the registers and instructions that make up the transactional memory (TM) facility. Transactional memory is a shared-memory synchronization construct allowing an application to perform a sequence of storage accesses that appear to occur atomically with respect to other threads.

A set of instructions, special-purpose registers, and state bits in the MSR (see Book III) are used to control a transactional facility that is associated with each hardware thread. A tbegin. instruction is used to initiate transactional execution, and a tend. instruction is used to terminate transactional execution. Loads and stores that occur between the tbegin. and tend. instructions appear to occur atomically. An implementation may prematurely terminate transactional execution for a variety of reasons, rolling back all transactional storage updates that have been made by the thread since the tbegin. was executed, and rolling back the contents of a subset of the thread's Book I registers to their contents before the tbegin. was executed. In the event of such premature termination, control is transferred to a software failure handler associated with the transaction, which may then retry the transaction or choose an alternate path depending on the cause of transaction failure. A transaction can be explicitly aborted via a set of conditional abort instructions and an unconditional abort instruction, tabort.. A tsr. instruction is used to suspend or resume transactional execution, while allowing the transaction to remain active.

\section*{Programming Note}

A tbegin. should always be followed immediately by a beq as the first instruction of the failure handler, that branches to the main body of the failure handler. The failure handler should always either retry the transaction or use non-transactional code to perform the same operation. (The number of retries should be limited to avoid the possibility of an infinite loop. The limit could be based on the perceived permanence / transience of the failure.) A failure handler policy which includes trying a different transaction before returning to the one that failed may fail to make forward progress.

\section*{Programming Note}

In code that may be executed transactionally, conditional branches should hint in favor of successful transactional execution where such a distinction exists. For example, the branch immediately following tbegin. should hint that the branch is very likely not to be taken. As another example, consider a method of coding a failure handler that executes the body of a transaction non-transactionally by branching past the TM control instructions (e.g. tsuspend.). Branches that bypass the TM control instructions should also hint that the branch is very likely not to be taken. These predictions will improve the efficiency of transactional execution, and may also help prevent the addition of spurious accesses to the transactional footprint.

\section*{Programming Note}

The architecture does not include a "fairness guarantee" or a "forward progress" guarantee for transactions. If two processors repeatedly conflict with one another in an attempt to complete a transaction, one of the two may always succeed while the other may always fail. If two processors repeatedly conflict with one another in an attempt to complete a transaction, both may always fail, depending on the details of the transaction. This is different from the behavior of a typical locking routine, in which one or the other of the competitors will generally get the lock.

Transactions performed using this facility are "strongly atomic", meaning that they appear atomic with respect to both transactional and non-transactional accesses performed by other threads. Transactions are isolated from reads and writes performed by other threads; i.e., transactional reads and writes will not appear to be interleaved with the reads and writes of other threads.

Nesting of transactions is supported using a form of nesting called "flattened nesting," in which transactions that are initiated during transactional execution are subsumed by the pre-existing transaction. Consequently, the effects of a nested transaction do not become visible until the outer transaction commits, and if a nested transaction fails, the entire set of transactions (outer as well as nested) is rolled back, and control is transferred to the outer transaction's failure handler. The memory barriers created by tbegin. and tend. and the integrated cumulative memory barrier that are described in Section 1.8, "Transactions" are only created for outer transactions and not for any transactions nested within them.

References to Store instructions, and stores, include dcbz and the storage accesses that it causes.

\section*{Rollback-Only Transactions}

Rollback-Only Transactions (ROTs) differ from normal transactions in that they are speculative but not atomic. They are initated by a unique variant of tbegin. They may be nested with other ROTs or with normal transactions. When a normal transaction is nested within a ROT, the behavior from the normal tbegin. until the end of the outer transaction is characteristic of a normal transaction. Although subject to failure from storage conflicts, the typical cause of ROT failure is via a Tabort variant that is executed after the program detects an error in its (software) speculation. Except where specifically differentiated or where differences follow from specific differentiation, the following description applies to ROTs as well as normal transactions.

\subsection*{5.1.1 Definitions}

Commit: A transaction is said to commit when it successfully completes execution. When a transaction is committed, its transactional accesses become irrevocable, and are made visible to other threads. A transaction completes by either commiting or failing.
Checkpointed registers: The set of registers that are saved to the "checkpoint area" when a transaction is initiated, and restored upon transaction failure, is a subset of the architected register state, consisting of the General Purpose Registers, Floating-Point Registers, Vector Registers, Vector-Scalar Registers, and the following Special Registers and fields: CR fields other than CRO, LR, CTR, FPSCR, AMR, PPR, VRSAVE, VSCR, DSCR, and TAR. The checkpointed registers include all problem state writable registers with the exception of CRO, LMRR, LMSER, EBBHR, EBBRR, BESCR, the Performance Monitor registers, and the Transactional Memory registers. With the exception of updates of CRO, and the Transactional Memory registers, explicit updates of registers that are not included in the set of checkpointed registers are disallowed in Transactional state (i.e., will cause the transaction to fail), but are permitted in Suspended state. Suspended state modifications of these registers will not be rolled back in the event of transaction failure. (Modifications of Transactional Memory registers are permitted in Non-transactional state, and modifications of the TFHAR are also permitted in Suspended state. Other attempts to modify Transactional Memory registers will cause a TM Bad Thing type Program interrupt.)

\section*{Programming Note}

CRO, and the Transactional Memory registers (TFHAR, TEXASR, TFIAR) are not saved and are not restored when the transaction fails because restoring them would lose information needed by the failure handler. The Performance Monitor registers, the event-based branching registers (BESCR, EBBHR, EBBRR), and the Load Monitored registers (LMRR, LMSER) are not saved or restored because saving and restoring them would add significant implementation complexity and is not needed by software. Also, these registers, except EBBHR, LMRR, and LMSER can be modified asynchronously by the processor, so restoring them when the transaction fails could cause loss of information.

Transactional accesses: Data accesses that are caused by an instruction that is executed when the thread is in the Transactional state (see Section 5.2) are said to be "transactional," or to have been "performed transactionally." The set of accesses caused by a committed normal transaction is performed as if it were a single atomic access. That is, it is always performed in its entirety with no visible fragmentation. The sets performed by normal transactions are thus serial-
ized: each happens in its entirety in some order, even when that order is not specified in the program or enforced between processors. Until a transaction commits, its set of transactional accesses is provisional, and will be discarded should the transaction fail. The set of transactional accesses is also referred to as the "transactional footprint."

Non-transactional accesses: Storage accesses performed in the existing Power storage model are said to be "non-transactional." In contrast to transactional storage accesses, there is no provision of atomicity across multiple non-transactional accesses. Non-transactional storage updates are not discarded in the event of a transaction failure.

Outer transaction: A transaction that is initiated from the Non-transactional state is said to be an outer transaction. A tbegin. instruction that initiates an outer transaction is sometimes referred to as an "outer tbegin.." Similarly, a tend. instruction with \(A=0\) that ends an outer transaction is sometimes referred to as an "outer tend."

Nested Transaction: A transaction that is initiated while already executing a transaction is said to be "nested" within the pre-existing transaction. The set of active nested transactions forms a stack growing from the outer transaction. A tend. with \(\mathrm{A}=0\) will remove the most recently nested transaction from the stack.

Failure: A transaction failure is an exceptional condition causing the transactional footprint to be discarded, the checkpointed registers to be reverted to their pre-transactional values, and the failure handler to get control.
Failure handler: A failure handler is a software component responsible for handling transaction failure. On transaction failure, hardware redirects control to the failure handler associated with the outer transaction.

Conflict: A transactional storage access is said to conflict with another transactional or non-transactional storage access if the two accesses overlap-i.e. if there is at least one byte that is referenced by both accesses-and at least one of the accesses is a store. If two transactions make conflicting accesses, at least one of them will fail. If a transaction fails as a result of a conflict with a store, the store may have been executed by another processor or may have been executed in Suspended state by the processor with the failing transaction. For a ROT, no conflict is caused if the ROT performs a load and another program performs a non-transactional store to the same storage location. The granularity at which conflict between storage accesses is detected is implementation-dependent, and may vary between accesses, but is never larger than a cache block.

A transactional storage access is said to conflict with a tlbie or slbieg if the storage location being accessed is in the page or segment the translation for which is
| being invalidated by the tlbie or slbieg. For a ROT, no conflict is caused if the access is a load.

A Suspended state cache control instruction is said to cause a conflict if it would cause the destruction of a transactional update or if it would make a transactional update visible to another thread.

\section*{Programming Note}

Warning: In descriptions of the transactional memory facility that precede V. 2.07B, the granularity at which conflict between storage accesses is detected was specified to be the cache block. Programs that were based on these early descriptions and depend on this granularity may need to be revised so as not to depend on it.
A future version of the architecture may define "transaction conflict granule", as the aligned unit of storage having the property that the granularity at which conflict between storage accesses is detected is never larger than the transaction conflict granule. The size of the transaction conflict granule would be implementation-dependent and would be added to the list of parameters useful to application programs in Section 4.1 and the last sentence of the first paragraph of the definition of "conflict" would use "transaction conflict granule" instead of "cache block".

\subsection*{5.2 Transactional Memory Facility States}

The transactional memory facility supports several modes of operation, referred to in this document as the "transaction state." These states control the behavior of storage accesses made during the transaction and the handling of transaction failure. Changes to transaction state affect all transactions currently using the transactional facility on the affected thread: the outer transaction as well as any nested transactions, should they exist.

Non-transactional: The default, initial state of execution; no transaction is executing. The transactional facility is available for the initiation of a new transaction.

Transactional: This state is initiated by the execution of a tbegin. instruction in the Non-transactional state. Storage accesses (data accesses) caused by instructions executed in the Transactional state are performed transactionally. Other storage accesses associated with instructions executed in the Transactional state (instruction fetches, implicit accesses) are performed non-transactionally. In the event of transaction failure, failure is recorded as defined in Section 5.3.2, and control is transferred to the failure handler as described in Section 5.3.3.

Suspended: The Suspended execution state is explicitly entered with the execution of a tsuspend. form of \(\boldsymbol{t s r}\). instruction during a transaction, the execution of a trechkpt. instruction from Non-transactional state, or as a side-effect of an interrupt while in the Transactional state. Storage accesses and accesses to SPRs that are not part of the checkpointed registers are performed non-transactionally; they will be performed independently of the outcome of the transaction. The initiation of a new transaction is prevented in this state. In the event of transaction failure, failure recording is performed as defined in Section 5.3.2, but failure handling is usually deferred until transactional execution is resumed (see Section 5.3.3 for details).
Until failure occurs, Load instructions that access storage locations that were transactionally written by the same thread will return the transactionally written data. After failure is detected, but before failure handling is performed, such loads may return either the transactionally written data, or the current non-transactional contents of the accessed location. The tcheck instruction can be used to determine whether any previous such loads may have returned non-transactional contents.

Suspended state Store instructions that access storage locations that have been accessed transactionally (due to load or store) by the same thread cause the transaction to fail.

\section*{Programming Note}

The intent of the Suspended execution state is to temporarily escape from transactional handling when transactional semantics are undesirable. Examples of such cases include storage updates that should be retained in the event of transactional failure, which is useful for debugging, interthread communication, the access of Caching Inhibited storage, and the handling of interrupts. In the event of transaction failure during the Suspended execution state, failure handling is deferred until transactional execution is resumed, allowing the block of Suspended state code to complete its activities.

\section*{Programming Note \\ During Suspended state execution, accessing storage locations that have been transactionally accessed by the same thread prior to entering Suspended state requires special care, because failure may occur due to uncontrollable events such as interactions with other threads or the operating system. Up until a transaction fails, loads from transactionally modified storage locations will return the transactionally modified data. However once the transaction fails, the loads may return either the transactionally updated version of storage, or a non-transactional version. Suspended state stores to transactionally modified blocks cause the thread's transaction to fail.}

Table 9 enumerates the set of Transactional Memory instructions and events that can cause changes to the transaction state. Transaction states are abbreviated N (Non-transactional), T (Transactional), and S (Sus| pended). (Interrupts, and the rfebb, rfid, rfscv, hrfid, and mtmsrd instructions, can also cause changes to the transaction state; see Book III.)

\section*{Programming Note}
tbegin. in Suspended state merely updates CRO. When tbegin. is followed by beq, this will result in a transfer to the failure handler. Nothing more severe (e.g. an interrupt) is required. The failure handler for a transaction for which initiation may be attempted in Suspended state should test CRO to determine whether tbegin. was executed in Suspended state. If so, it should attempt to emulate the transaction non-transactionally. (This case can arise, for example, if a transaction enters Suspended state and then calls a library routine that independently attempts to use transactions.)

Notice that, although a failure handler runs in Non-transactional state when reached because the transaction has failed, it runs in Suspended state for the case discussed in this Programmng Note.)
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline \begin{tabular}{l}
Instr/ Event \\
State
\end{tabular} & tbegin. & tend. & Abort caused by tabort. and conditional tabort. variants & tsuspend. & tresume. & Failure & treclaim. & trechkpt. \\
\hline N & T & \(\mathrm{N}^{2}\) & \(\mathrm{N}^{2}\) & \(N^{2}\) & \(N^{2}\) & Not applicable & \(N^{6}\) & \(S^{7}\) \\
\hline T & T & N , if outer transaction or \(\mathrm{A}=1\) form; otherwise T & \(\mathrm{N}^{3,4}\) & S & T & \(\mathrm{N}^{3,4}\) & \(N^{3}\) & \(S^{6}\) \\
\hline S & \(\mathrm{S}^{1}\) & \(S^{6}\) & \(s^{3}\) & \(s^{2}\) & \(T^{5}\) & \(s^{3}\) & \(\mathrm{N}^{3}\) & \(S^{6}\) \\
\hline
\end{tabular}

\section*{Notes}
1. CRO updated indicating transactional initiation was unsuccessful, due to a pre-existing transaction occupying the transactional facility.
2. Execution of these operations does not affect transaction state, allowing for the instructions to be used in software modules called from Non-transactional, Transactional, and Suspended paths.
3. If failure recording has not previously occurred, failure recording is performed as defined in Section 5.3.2.
4. Failure handling is performed as defined in Section 5.3.3.
5. If failure has occurred during Suspended execution, failure handling will be performed sometime after the execution of tresume, and no later than the set of events listed in Section 5.3.3.
6. Generate TM Bad Thing type Program interrupt.
7. If TEXASR \(_{\mathrm{FS}}=0\), generate a TM Bad Thing type Program interrupt.

Table 9: Transaction state transitions caused by TM instructions and transaction failure

\subsection*{5.2.1 The TDOOMED Bit}

The status of an active transaction is summarized by a transaction doomed bit (TDOOMED) that resides in an implementation-dependent location. When 0 , it indicates that the active transaction is valid, meaning that it remains possible for the transaction to commit successfully, if failure does not occur before committing. When 1 it indicates that transaction failure has already occurred for the transaction.

The TDOOMED bit is set to 0 upon the successful initiation of an outer transaction by tbegin.. It is set to 1 when failure occurs or as a result of executing trechkpt.. When failure occurs, TDOOMED is set to 1 before any other effects of the transaction failure (recording the failure in TEXASR, rollback of transactional stores, over-writing of the transactionally accessed locations by a conflicting store, etc.) are visible to software executing on the processor that executed the transaction. In Non-transactional state, the value of TDOOMED is undefined.

\subsection*{5.3 Transaction Failure}

\subsection*{5.3.1 Causes of Transaction Failure}

A transaction failure is said to be "externally-induced" if the failure is caused by a thread other than the transactional thread. Likewise, a transaction failure is said to be "self-induced" if the failure is caused by the transactional thread itself. In the list of failure causes below, for self-induced failures that result from copy or paste[.] instructions, if the reference goes through CSM and addresses (local) main storage, it will be reported as an externally induced failure instead of a self-induced failure.
For self-induced failure as a result of attempting to execute an instruction that is forbidden in the Transactional state, a Privileged Instruction type of Program Interrupt takes precedence over transaction failure. (For example, an attempt to execute stdcix in Transactional state and problem state will result in a Privileged Instruction type of Program interrupt.) Transaction failure takes precedence over all other interrupt types. The relevant instructions are listed in the fourth bullet of the second set of bullets below and the first bullet in the third set of bullets below.

In general, a ROT will not fail in the following scenarios when the failure is specified as a conflict on a transactional access and the access is a load.

Transactions will fail for the following externally-induced causes
- Conflict with transactional access by another thread
- Conflict with non-transactional access by another thread
■ In either of the previous two cases, if a successful Store Conditional would have conflicted, but the Store Conditional is not successful, it is implemen-tation-dependent whether a conflict is detected
- Conflict with a translation invalidation caused by a tlbie or slbieg performed by another thread
- paste[.] to a block that was previously accessed transactionally is executed by another thread.
- copy from a block that was previously written transactionally is executed by another thread.
Transactions will fail for the following self-induced causes
- Termination caused by the execution of tabort., tabortdc., tabortdci., tabortwc., tabortwci. or treclaim. instruction.
- Transaction level overflow, defined as an attempt to execute tbegin. when the transaction level is already at its maximum value
- Footprint overflow, defined as an attempt to perform a storage access in Transactional state which exceeds the capacity for tracking transactional accesses.
- Execution of the following instructions while in the Transactional state: icbi, copy, paste[.], cp_abort, Iwat, Idat, stwat, stdat, dcbf, dcbi, dcbst, rfscv, [h]rfid, rfebb, mtmsr[d], msgsnd, msgsndp, msgclr, msgclrp, slbie[g], slbia, slbmte, slbfee, stop, and tlbie[I]. (These instructions are considered to be disallowed in Transactional state.) The disallowed instruction is not executed; failure handling occurs before it has been executed.

\section*{- Programming Note}

Note that execution of a stop instruction in Suspended state causes a TM Bad Thing type Program interrupt.
- Execution, while in Transactional state, of mtspr specifying an SPR that is not part of the checkpointed registers and is not a Transactional Memory SPR. The mtspr is not executed; failure handling occurs before it has been executed. (Modification of XER \({ }_{\text {FXCC }}\) and \(\mathrm{CR}_{\mathrm{CRO}}\) are allowed, but the changes will not be rolled back in the event of transaction failure.)
- Conflict caused by a Suspended state store to a storage location that was previously accessed transactionally. If the store would have been performed by a successful Store Conditional instruction, but the Store Conditional instruction does not succeed, it is implementation-dependent whether a conflict is detected.

■ Conflict caused by a Suspended state Load Atomic or Store Atomic instruction updating a block that was previously accessed transactionally.
■ paste[.] to a block that was previously accessed transactionally is executed in Suspended state on the thread executing the transaction.
■ Conflict caused by a Suspended state tlbie or slbieg that specifies a translation that was previously used transactionally. (This case will be recorded as a translation invalidation conflict because it may be hard to differentiate from a conflict caused by a tlbie or slbieg performed by another thread and because it is highly likely to be a transient failure.)
For each of the following potential causes, the transaction will fail if the absence of failure would compromise transaction semantics; otherwise, whether the transaction fails is undefined.
- Execution of the following instructions while in the Transactional state: Ibzcix, Idcix, Ihzcix, Iwzcix, stbcix, stdcix, sthcix, stwci. The disallowed instruction is not executed; failure handling occurs before it has been executed. (These instructions are considered to be disallowed in Transactional state if they cause transaction failure in Transactional state.) Execution of these instructions in the Suspended state is allowed and does not cause transaction failure.
- Execution of the following instruction in the Transactional state: wait. The disallowed instruction is not executed; failure handling occurs before it has been executed. (This instruction is considered to be disallowed in a transaction if it causes transaction failure.)
- Execution of the following instruction in the Suspended state: wait. The disallowed instruction is treated as a no-op; failure recording occurs. (This instruction is considered to be disallowed in a transaction if it causes transaction failure.)
- Access of a disallowed type while in the Transactional state: Caching Inhibited, Write Through Required, and Memory Coherence not Required for data access; Caching Inhibited for instruction fetch. The disallowed access is not performed; failure handling occurs such that the instruction that would cause (or be associated with, for instruction fetch) the disallowed access type appears not to have been executed. Accesses of this type in the Suspended state are allowed and do not cause transaction failure.
- Instruction fetch from a storage location that was previously written transactionally (reported as a unique cause that includes both self-induced and externally-induced instances)

■ dcbf, dcbi, or icbi specifying a block that was previously accessed transactionally, in either of the following cases.

\section*{Programming Note}

Note that dcbf with L=3 should never compromise transactional semantics, but it is still permitted to cause transaction failure in Suspended state and it is disallowed in Transactional state.
- the instruction (dcbf, dcbi, or icbi) is executed in Suspended state on the processor executing the transaction (self-induced conflict)
- the instruction is executed by another processor (externally-induced conflict)
■ dcbst specifying a block that was previously written transactionally, in either of the following cases.
■ dcbst is executed in Suspended state on the processor executing the transaction (self-induced conflict)
■ dcbst is executed by another processor (externally-induced conflict)
- copy, in any of the following cases.
- copy from a block that was previously accessed transactionally is executed in Suspended state on the processor executing the transaction (self-induced conflict)
- copy from a block that was previously accessed transactionally is executed by another processor (externally-induced conflict)
- Cache eviction of a block that was previously accessed transactionally
Transactions may also fail due to implementation-specific characteristics of the transactional memory mechanism.

\section*{Programming Note}

Warning: Software should not depend for its correct execution on the behavior (whether or not the relevant transaction fails) of the cases described in the preceding set of bullets. The behavior is likely to vary from design to design. Such a dependence would impact the software's portability without any tangible advantage.

\section*{Programming Note}

Because the atomic nature of a transaction implies an apparent delay of its component accesses until they can be performed in unison, the use of cache control instructions to manage cache residency and/or the performing of storage accesses may have unexpected consequences. Although they may not cause transaction failure directly, their use in a transaction is strongly discouraged.

If an instruction or event does not cause transaction failure, it behaves as defined in the architecture.

The set of failure causes and events are further classified as "precise" and "imprecise" failure causes. All
externally induced events are imprecise, and all self-induced events are precise with the exception of the following cases:
- Self-induced conflicts caused by instruction fetch
- Self-induced conflicts caused by footprint overflow
- Self-induced conflicts in Suspended state (because failure handling is deferred in Suspended state).

When failure recording and handling occur (as defined in Section 5.3.2 and 5.3.3) for a precise failure, in general they will occur precisely according to the sequential execution model, adhering to the following rules: The only exception is that if an mtspr sequence started by mtgsr (see Section 4.4.5 of Book III) is active when the precise failure occurs, some of the sequence's mtsprs beyond the point at which the recording and handling occur may have been executed; see Chapter 11 of Book III. Statements elsewhere in this chapter that a given failure is precise do not preclude the mtspr case just described.
1. Effects of the failure occur such that all instructions preceding the instruction causing the failure appear to have completed with respect to the executing thread.
2. The instruction causing the failure may appear not to have begun execution (except for causing the failure), or may have completed, depending on the failure cause.
3. Architecturally, no subsequent instruction has begun execution.
Failure handling for imprecise failure types is guaranteed to occur no later than the execution of tend. with \(\mathrm{A}=1\) or \(\mathrm{TEXASR}_{\mathrm{TL}}=1\). Failure recording for imprecise failure types is guaranteed to occur no later than failure handling. Any operation that can cause imprecise failure if performed in-order can also cause imprecise failure if performed out-of-order.

\section*{Programming Note}

Because instruction fetch from a transactionally modified storage location may result in transaction failure, and because conflict between storage accesses may be detected at granularity as large as a cache block, it is recommended that instructions and transactionally accessed data not be co-located within a single cache block.

\section*{Programming Note}

The architecture does not detect and cause transaction failure for translation invalidations to transactionally accessed pages or segments, when the translation invalidation is caused by instructions other than tlbie or slbieg (i.e., slbie, slbia, tlbiel). Consequently, software is responsible for terminating transactions in circumstances where such local translation invalidations may affect a local transaction.

\subsection*{5.3.2 Recording of Transaction Failure}

When transaction failure occurs, information about the cause and circumstances of failure are recorded in SPRs associated with the transactional facility. Failure recording is performed a single time per transaction that fails, controlled by the state of the TEXASR failure summary (FS) bit; when 0, FS indicates that failure recording has not already been performed, and is therefore permissible.

The following RTL function specifies the actions taken during the recording of transaction failure:
```

TMRecordFailure(FailureCause)
\#FailureCause is 32-bit cause
code
if TEXASR $_{F S}=0$
if failure IA known then
TFIAR $\leftarrow$ CIA
TEXASR $_{37} \leftarrow 1$
else
TFIAR $\leftarrow$ approximate instruction address
$\mathrm{TEXASR}_{37} \leftarrow 0$
$\operatorname{TEXASR}_{0: 31} \leftarrow$ FailureCause
if MSR $_{T S}=0 \mathrm{~b} 01$ then $\mathrm{TEXASR}_{\text {Suspended }} \leftarrow 1$
TEXASR $_{\text {PRIVILEGE }} \leftarrow$ MSR $_{\text {HV }} \| \operatorname{MSR}_{\text {PR }}$
TFIAR $_{\text {PRIVILEGE }} \leftarrow$ MSR $_{\text {HV }} \| \operatorname{MSR}_{\text {PR }}$
TEXASR $_{\text {FS }} \leftarrow 1$
TDOOMED $\leftarrow 1$

```

When failure recording occurs, the TEXASR and TFIAR SPRs are set indicating the source of failure. When possible, TFIAR is set to the effective address of the instruction that caused the failure, and TEXASR 37 is set to 1 indicating that the contents of TFIAR are exact. When the instruction address is not known exactly, an approximate value is placed in TFIAR and TEXASR 37 is set to 0 . TEXASR bits \(0: 31\) are set indicating the cause of the failure, and the TEXASR Suspended TEXASR \(_{\text {Privilege }}\), and TFIAR Privilege fields are set indicating the machine state in which the failure was recorded. \(\mathrm{TEXASR}_{\mathrm{TL}}\) is unchanged. The TDOOMED bit is set to 1.

> Programming Note TFIAR is intended for use in the debugging of transactional programs by identifying the source of transaction failure. Because TFIAR may not always be set exactly, software should test TEXASR 37 before use; if zero, the contents of TFIAR are an approximation.

\subsection*{5.3.3 Handling of Transaction Failure}

Discarding of the transactional footprint may begin immediately after detection of failure and, except in the case of an abort in Suspended state, may continue until the rest of failure handling is complete. However, the timing of the rest of failure handling is dependent on the state of the transactional facility. In the case of an abort in Suspended state, the transactional footprint is discarded immediately, despite that the rest of failure handling is deferred.
In Transactional state, failure handling may occur immediately, but an implementation is free to delay handling until one of the following failure synchronizing events occurs in Transactional state.
- An abort caused by the execution of a tabort., tabortdc., tabortdci., tabortwc., or tabortwci. instruction.
- The execution of a treclaim. instruction.
- An attempt, in Transactional state, to execute a disallowed instruction, perform an access of a disallowed type, or execute an mtspr instruction that specifies an SPR that is not part of the checkpointed registers and is not a Transactional Memory SPR.
- Nesting level overflow.
- An attempt to transition from Transactional to Suspended state caused by tsuspend. or by an interrupt or event.
- An attempt to commit a transaction, caused by the execution of tend. with A \(=1\) or when \(\mathrm{TEXASR}_{\mathrm{TL}}=\) 1.

When a failure synchronizing event occurs in Transactional state, the processor waits until all preceding Transactional and Suspended state loads have been performed with respect to all processors and mechanisms and all failures that have occurred up to that point have been recorded. Then failure handling occurs if a failure has been recorded; otherwise, processing of the failure synchronizing event continues. If failure is caused by the failure synchronizing event, failure handling occurs immediately.

When failure handling occurs, checkpointed registers are reverted to their pre-transactional values, the transactional footprint is discarded if it has not previously been discarded, and any resources occupied by the transaction are discarded. If the failure is not caused by
treclaim., the following things occur. CRO is set to Ob101 II 0. The transaction state is set to Non-transactional, and control flow is redirected to the instruction address stored in TFHAR. If the failure is caused by treclaim., CRO is not set to indicate failure and the transaction's failure handler is not invoked.

The following RTL function specifies the actions taken during the handling of transaction failure:
```

TMHandleFailure()
If the transactional footprint has not previ-
ously been discarded
Discard transactional footprint
Revert checkpointed registers to pre-transac-
tional values
Discard all resources related to current trans-
action
MSR TS }\leftarrow0\mathrm{ Ob00 \#Non-transactional
If failure was not caused by treclaim.,
NIA \leftarrow TFHAR
CR0}\leftarrow0\textrm{b}101|

```

Upon failure detected in Suspended state from causes other than the execution of a treclaim. instruction, failure handling is deferred until the transaction is resumed. Once resumed, failure handling will occur no later than the set of failure synchronizing events listed above. Upon failure in Suspended state caused by treclaim., failure handling is immediate (but CRO is not set to indicate failure and the transaction's failure handler is not invoked).

\section*{Programming Note}

A Load instruction executed immediately after treclaim. or a conditional or unconditional Abort instruction is guaranteed not to load a transactional storage update.

\subsection*{5.4 Transactional Memory Facility Registers}

The architecture is augmented with three Special Purpose Registers in support of transactional memory. TFHAR stores the effective address of the software failure handler used in the event of transaction failure. TFIAR is used to inform software of the exact location of the transaction failure, when possible. TEXASR contains a transaction level indicating the nesting depth of an active transaction, as well as an indicator of the cause of transaction failure and some machine state when the transaction failed. These registers can be written only when in Non-transactional state, and for TFHAR, also when in Suspended state.

\subsection*{5.4.1 Transaction Failure Handler Address Register (TFHAR)}

The Transaction Failure Handler Address Register is a 64-bit SPR that records the effective address of a software failure handler used in the event of transaction failure. Bits 62:63 are reserved.


Figure 5. Transaction Failure Handler Address Register (TFHAR)

This register is written with the NIA for the tbegin. as a side-effect of the execution of an outer tbegin. instruction (tbegin. executed in the Non-transactional state).

\subsection*{5.4.2 Transaction EXception And Status Register (TEXASR)}

The Transaction EXception And Status Register is a 64-bit register, containing a transaction level (TEXASR \({ }_{\text {TL }}\) ) and status information for use by transaction failure handlers. The identification of the cause and persistence of transaction failure reported in bits 7:30 may rarely be inaccurate. Bits 0:31 are called the failure cause in the instruction descriptions.


Figure 6. Transaction EXception And Status Register (TEXASR)


Figure 7. Transaction EXception And Status Register Upper (TEXASRU)

\section*{Bit(s Description}

0:6 Failure Code
The Failure Code is copied from the tabort. or treclaim. source operand. When set, TFIAR is exact.

\section*{\(7 \quad\) Failure Persistent}

The failure is likely to recur on each execution of the transaction. This bit is set to 1 for causes in bits 8:11, copied from the tabort. or treclaim. source operand when RA is nonzero, and set to 0 for all other failure causes.

\section*{Programming Note}

The Failure Persistent bit may be viewed as an eighth bit in the failure code in that both fields are supplied by the least significant byte of RA and software may use all eight to differentiate among the cases for which it performs an abort or reclaim. However, software is expected to organize its cases so that bit 7 predicts the persistence of the case.

\section*{Programming Note}

Warning: Software must not depend on the value of the Failure Persistent bit for correct execution. The number of retries for a transient failure should be counted, and a limit set after which the program will perform the operation non-transactionally. In the analysis of failures, consideration should be given to the fact that speculative execution can cause unexpected behavior.

The inaccuracy of the Failure Persistent bit arises from two causes. First, a kind of failure that is usually transient, such as conflict with another thread, may in certain unusual circumstances be persistent. Second, if the cause of transaction failure is identified incorrectly, the Failure Persistent bit will inherit this inaccuracy -- i.e., will be set to 0 or 1 based on the identified failure cause.

\section*{Disallowed}

The instruction, SPR, or access type is not permitted. When set, TFIAR is exact. See Section 5.3.1, "Causes of Transaction Failure".

\section*{Programming Note}

An instruction fetch to storage that is Caching Inhibited, while nominally disallowed, will be reported as Implementa-tion-specific (bit 15). This choice was made because it seems like a relatively unlikely programming error, and there is a significant chance that data from an external conflict (store by another thread) could indirectly cause a wild branch to storage that is Caching Inhibited.

\section*{Nesting Overflow}

The maximum transaction level was exceeded. When set, TFIAR is exact.

\section*{Footprint Overflow}

The tracking limit for transactional storage accesses was exceeded. When set, TFIAR is an approximation.

\section*{Programming Note}

Note that transactional footprint tracking resources may be shared by multiple programs executing concurrently. Depending on the circumstances, this failure cause may or may not be persistent.

11 Self-Induced Conflict
A self-induced conflict occurred in Suspended state, due to one of the following: a store to a storage location that was previously accessed transactionally; a dcbf, dcbi, or icbi specifying a block that was previously accessed transactionally; a dcbst specifying a block that was previously written transactionally; or a copy from or paste[.] to a block that was previously accessed transactionally. When set, TFIAR may be exact.
12 Non-Transactional Conflict
A conflict occurred with a non-transactional access by another processor. When set, TFIAR is an approximation.

13 Transaction Conflict
A conflict occurred with another transaction. When set, TFIAR may be exact.
14 Translation Invalidation Conflict
| A conflict occurred with a TLB or SLB invalidation. When set, TFIAR is an approximation.

15 Implementation-specific
An implementation-specific condition caused the transaction to fail. Such conditions are transient and the value in the TFIAR may be exact.

\section*{Instruction Fetch Conflict}

An instruction fetch (by this or another thread) was performed from a storage location that was previously written transactionally. Such conditions are transient and the value in the TFIAR may be exact.

17:30 Reserved for future failure causes

31 Abort
Termination was caused by the execution of a tabort., tabortdc., tabortdci., tabortwc., tabortwci. or treclaim. instruction. When due to tabort. or treclaim., bits in TEXASR \({ }_{0: 7}\) are user-supplied. When set, TFIAR is exact.

\section*{32 Suspended}

When set to 1, the failure was recorded in Suspended state. When set to 0 , the failure was recorded in Transactional state.

33 Reserved
34:35 Privilege
The thread was in this privilege state when the failure was recorded. This was the value
\(M S R_{H V}\) II \(M S R_{P R}\) when the failure was recorded.
Failure Summary (FS)
Set to 1 when a failure has been detected and failure recording has been performed.

\section*{37 TFIAR Exact}

Set to 1 when the value in the TFIAR is exact. Otherwise the value in the TFIAR is approximate.
ROT
Set to 1 when a ROT is initiated. Set to zero when a non-ROT tbegin. is executed.
Reserved
40:51 Reserved
52:63 Transaction Level (TL)
Transaction level (nesting depth +1 ) for the active transaction, if any; otherwise 0 if the most recently executed transaction completed successfully, or the transaction level at which the most recently executed transaction failed if the most recently executed transaction did not complete successfully.

\section*{Programming Note}

A value of 1 corresponds to an outer transaction. A value greater than 1 corresponds to a nested transaction.

The transaction level in TEXASR \(_{T L}\) contains an unsigned integer indicating whether the current transaction is an outer transaction, or is nested, and if nested, its depth. The maximum transaction level supported by a given implementation is of the form \(2^{t}-1\). The value of \(t\) corresponding to the smallest maximum is 4 ; the value of \(t\) corresponding to the largest maximum is 12 . This value is tied to the "Maximum transaction level" parameter useful for application programmers, as specified in Section 4.1. The high-order \(12-t\) bits of TEXASR \(_{T L}\) are treated as reserved.

Transaction failure information is contained in TEXASR \(_{0: 37}\). The fields of TEXASR are initialized upon the successful initiation of a transaction from the Non-transactional state, by setting TEXASR TL \(^{\text {to } 1, ~}\) indicating an outer transaction, and all other fields to 0 .
When transaction failure is recorded, the failure summary bit \(\operatorname{TEXASR}_{\text {FS }}\) is set to 1 , indicating that failure has been detected for the active transaction and that failure recording has been performed. TEXASR \({ }_{0: 31}\) are set indicating the source of the failure. Exactly one of bits 8 through 31 will be set indicating the instruction or event that caused failure. In the event of failure due to the execution of a tabort., tabortdc., tabortdci., tabortwc., tabortwci. or treclaim. instruction, TEXASR \(_{31}\) is set to 1 , and, for tabort. and treclaim., a software defined failure code is copied from a register
operand to TEXASR \(_{0: 7}\). TEXASR Suspended indicates whether the transaction was in the Suspended state at the time that failure occurred. The values of \(\mathrm{MSR}_{H V}\) and \(M S R_{P R}\) at the time that failure occurs are copied to TEXASR \(_{34}\) and TEXASR 35 , respectively. In some circumstances, the failure causing instruction address in TFIAR may not be exact. In such circumstances, TEXASR \(_{37}\) is set to 0 indicating that the contents of TFIAR are not exact; otherwise TEXASR \(3_{77}\) is set to 1 .

\section*{Programming Note}

The transaction level contained in TEXASR \({ }_{T L}\) should be interpreted by software as follows:
When in the Transactional or Suspended state, this field contains an unsigned integer representing the transaction level of the active transaction, with 1 indicating an outer transaction, and a number greater than 1 indicating a nested transaction. The nesting depth of the active transaction is TEXASR \(_{T L}-1\).
When in the Non-transactional state, TEXASR \(_{\text {TL }}\) contains 0 if the last transaction committed successfully, otherwise it contains the transaction level at which the most recent transaction failed.

\section*{Programming Note}

The Privilege bits in TEXASR represent the state of the machine at the point when failure occurs. This information may be used by problem state software to determine whether an unexpected hypervisor or operating system interaction was responsible for transaction failure. This information may be useful to operating systems or hypervisors when restoring register state for failure handling after the transactional facility was reclaimed, to determine which of the operating system or the hypervisor has retained the pre-transactional version of the checkpointed registers.

\subsection*{5.4.3 Transaction Failure Instruction Address Register (TFIAR)}

The Transaction Failure Instruction Address Register is a 64-bit SPR that is set to the exact effective address of the instruction causing the failure, when possible. Bits 62:63 contain the privilege state when the failure was recorded. This was the value \(M S R_{H V} \| M S R_{P R}\) when the failure was recorded.


Figure 8. Transaction Failure Instruction Address Register (TFIAR)
In certain cases, the exact address may not be available, and therefore TFIAR will be an approximation. An

\section*{Version 3.0}
approximate value will point to an instruction near the instruction that was executing at the time of the failure. TFIAR accuracy is recorded in an Exact bit residing in TEXASR 37 .

\subsection*{5.5 Transactional Facility Instructions}

Similar to the Floating-Point Status and Control Register instructions, modifications of transaction state caused by the execution of Transactional Memory instructions or by failure handling synchronize the effects of exception-causing floating-point instructions executed by a given processor. Executing a Transactinal Memory instruction, or invocation of the failure handler, ensures that all floating-point instructions previously initiated by the given processor have completed before the transaction state is modified, and that no subsequent floating-point instructions are initiated
by the given processor until the transaction state has been modified. In particular:
■ All exceptions that will be caused by the previously initiated instructions are recorded in the FPSCR before the transaction state is modified.
- All invocations of the system floating-point enabled exception error handler that will be caused by the previously initiated instructions have occurred before the transaction state is modified.
- No subsequent floating-point instruction that alters the settings of any FPSCR bits is initiated until the transaction state has been modified.
(Floating-point Storage Access instructions are not affected.)

\section*{Transaction Begin \\ X-form}
tbegin. \(\quad R\)
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 31 & 6 & 7 & & 16 & 21 & 654 & 1
31 \\
\hline
\end{tabular}
```

ROT \leftarrow R
CRO \leftarrow0 || MSR TS | | 0
if MSR TS = 0b00 then \#Non-transactional
TEXASR \leftarrow 0x000000000 || 0b00 || ROT || 0b0 ||
0x000001
TFHAR \leftarrow CIA + 4
TDOOMED \leftarrow0
MSR TS }\leftarrow0\textrm{Ob}1
checkpoint area \leftarrow (checkpointed registers)
if not ROT and the transaction succeeds then
else if MSR TS = 0b10 then \#Transactional
if TEXASR
cause \leftarrow 0x01400000
TMRecordFailure(cause)
TMHandleFailure()
else
TEXASR TL }\leftarrow\mp@subsup{T}{TEXASR }{TL
if (TEXASR ROOT=1) \& (not ROT)
\& the transaction succeeds
insert tbegin memory barrier
TEXASR ROT }\leftarrow

```
I

The tbegin. instruction initiates execution of a transaction, either an outer transaction or a nested transaction, as described below.

An outer transaction is initiated when tbegin. is executed in the Non-transactional state. If \(\mathrm{R}=0\) and the transaction is successful, the tbegin memory barrier described in Section 1.8 is inserted. TEXASR and TFHAR are initialized, and the TDOOMED bit is set to 0 . A nested transaction is initiated when tbegin. is executed in the Transactional state unless the transaction level is already at its maximum value, in which case failure recording is performed with a failure cause of
\(0 \times 01400000\) and failure handling is performed. When initiating a nested transaction, the transaction level held in TEXASR \({ }_{T L}\) is incremented by 1 , and if TEXASR ROT \(=1\) but \(\mathrm{R}=0\), and the transaction succeeds, the tbegin memory barrier described in Section 1.8 is inserted and TEXASR ROT is turned off. The effects of a nested transaction will not be visible until the outer transaction commits, and in the event of failure, the checkpointed registers are reverted to the pre-transactional values of the outer transaction. Initiation of a transaction is unsuccessful when in the Suspended state.

When successfully initiated, transactional execution continues until the transaction is terminated using a tend., tabort., tabortdc., tabortdci., tabortwc., tabortwci., or treclaim. instruction, suspended using a \(\boldsymbol{t s r}\) instruction, or failure occurs. Upon transaction failure while in the Transactional state, transaction failure recording and failure handling are performed as defined in Section 5.3. Upon transaction failure while in the Suspended state, failure recording is performed as defined in Section 5.3.2, but failure handling is usually deferred.

CRO is set as follows.
\begin{tabular}{l|l}
\hline CRO & Description \\
\hline \(000 \| 0\) & \begin{tabular}{c} 
Transaction initiation successful, \\
unnested (Transaction state of \\
Non-transactional prior to tbegin.)
\end{tabular} \\
\(0010 \| 0\) & \begin{tabular}{l} 
Transaction initiation successful, nested \\
(Transaction state of Transactional \\
prior to tbegin.)
\end{tabular} \\
\begin{tabular}{l} 
Transaction initiation unsuccessful, \\
(Transaction state of Suspended prior \\
to tbegin.)
\end{tabular}
\end{tabular}

Other than the setting of CRO, tbegin. in the Suspended state is treated as a no-op.

The use of the A field is implementation specific.

\section*{Special Registers Altered:} CRO TEXASR TFHAR TS

\section*{Programming Note}

When a transaction is successfully initiated, and failure subsequently occurs, control flow will be redirected to the instruction following the tbegin. instruction. When failure handling occurs, as described in Section 5.3.3, CR0 is set to Ob101 II 0. Consequently, instructions following tbegin. should also expect this value as an indication of transaction failure. Most applications will follow tbegin. with a conditional branch predicated on \(\mathrm{CRO}_{2}\); code at this target is responsible for handling the transaction failure.

\section*{Transaction End}

X-form
tend. A
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline & & & \(11 / /\) & & & \\
\hline \multicolumn{7}{|l|}{\multirow[t]{3}{*}{```
CR0}\leftarrow0\textrm{bb}||\mp@subsup{\textrm{MSR}}{\textrm{TS}}{}||
if MSR TS = 0b10 then #Transactional
    if A = 1 |EXASR TTL = 1 then
        if (TDOOMED) then
            TMHandleFailure()
        else
            if not TEXASR ROT
                insert integrated cumulative
                memory barrier
            Commit transaction
            TEXASR TL }\leftarrow
                Discard all resources related to current
transaction
            MSR
            if not TEXASR ROT
                insert tend memory barrier
    else TEXASR 
```}} \\
\hline & & & & & & \\
\hline & & & & & & \\
\hline
\end{tabular}

The \(A=0\) variant of tend. supports nested transactions, in which the transaction is committed only if the execution of tend. completes an outer transaction. Execution of this variant by a nested transaction (TEXASR \({ }_{T L}>1\) ) causes \(\operatorname{TEXASR}_{\text {TL }}\) to be decremented by 1 . The \(A=1\) variant of tend. unconditionally completes the current outer transaction and all nested transactions.

When the tend. instruction completes an outer transaction, transaction commit is predicated on the TDOOMED bit. If TDOOMED is 1 , failure handling occurs as defined in Section 5.3.3. If TDOOMED is 0 , the transaction is committed, and TEXASR \({ }_{\text {TL }}\) is set to 0 . In both cases, the transaction state is set to Non-transactional.

When the tend. instruction commits a transaction, it atomically commits its writes to storage. If TEXASRROT \(=0\), the integrated cumulative memory barrier is inserted prior to the creation of the aggregate store, and the tend memory barrier described in Section 1.8 is inserted after the aggregate store. If the transaction has failed prior to the execution of tend., no storage updates are performed and no memory barrier is inserted. In either case (success or failure), all resources associated with the transaction are discarded.

If the transaction succeeds, Condition Register field 0 is set to \(0 \| \mathrm{MSR}_{\text {TS }}\) II 0 . If the transaction fails, CRO is set to Ob101 II 0.

Other than the setting of CRO, tend. in Non-transactional state is treated as a no-op. If an attempt is made to execute tend. in Suspended state, a TM Bad Thing type Program interrupt occurs.

\section*{Special Registers Altered:}

\section*{CR0 TEXASR TS}

\section*{Extended Mnemonics}

Examples of extended mnemonics for Transaction End.

\section*{Extended: Equivalent To: \\ tend. \\ tend. 0}
tendall.

\section*{Programming Note}

When an outer tend. or a tend. with \(A=1\) is executed in the Transactional state, the CRO value Ob101 II 0 will never be visible to the instruction that immediately follows tend., because in the event of failure the failure handler will have been invoked not later than the completion of the tend. instruction.

Transaction Abort X-form
tabort. RA
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{1}{|l|}{ I/I } & RA & I/I & & 910 & 1 \\
0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

CRO}\leftarrow0|M\mp@subsup{MSR}{TS}{}||
if MSR
\#Transactional, or Suspended
if RA = 0 then cause \leftarrow0x00000001
else cause \leftarrowGPR(RA)56:63 || 0x000001
if MSR
Discard the transactional footprint
TMRecordFailure(cause)
if MSR TSS = 0b10 then
\#Transactional
TMHandleFailure()

```

The tabort. instruction sets condition register field 0 to 0 II \(\mathrm{MSR}_{\text {TS }}\) II 0 . When in the Transactional state or the Suspended state the tabort. instruction causes transaction failure, resulting in the following:

Failure recording is performed as defined in Section 5.3.2. If RA is 0 , the failure cause is set to \(0 \times 00000001\), otherwise it is set to \(\operatorname{GPR}(\mathrm{RA})_{56: 63}\) II \(0 \times 000001\).
If the transaction state is Transactional, failure handling is performed as defined in Section 5.3.3 (this includes discarding the transactional footprint).
If the transaction state is Suspended, the transactional footprint is discarded (if not already discarded for a pending failure), but failure handling is deferred.
Other than the setting of CRO, execution of tabort. in the Non-transactional state is treated as a no-op.

Special Registers Altered:
CRO TEXASR TFIAR TS

\section*{Transaction Abort Word Conditional X-form}
tabortwc. TO,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 31 & TO & RA & RB & & 782 & 1 \\
0 & & 6 & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

a\leftarrow EXTS((RA) 32:63)
b}\leftarrow\operatorname{EXTS}((\textrm{RB}\mp@subsup{)}{32:63}{\prime}
abort }\leftarrow
CRO \leftarrow 0 || MSR TS || 0
if (a<b) \& TOO then abort }\leftarrow
if (a > b) \& TO
if (a=b) \& TOO2 then abort }\leftarrow
if (a u< b) \& TO
if (a >u b) \& TO
if abort \& (MSR
\#Transactional or Suspended
cause }\leftarrow0x0000000
if MSR RS = 0.b01 \& TEXASR
Discard transactional footprint
TMRecordFailure(cause)
if MSR
TMHandleFailure()

```

The tabortwc. instruction sets condition register field 0 to 0 II \(\mathrm{MSR}_{\mathrm{TS}}\) II 0 . The contents of register \(\mathrm{RA}_{32: 63}\) are compared with the contents of register \(\mathrm{RB}_{32: 63}\). If any bit in the TO field is set to 1 and its corresponding condition is met by the result of the comparison, and the transaction state is Transactional or Suspended, then the tabortwc. instruction causes transaction failure, resulting in the following:
Failure recording is performed as defined in Section 5.3.2, using the failure cause \(0 \times 00000001\).

If the transaction state is Transactional, failure handling is performed as defined in Section 5.3.3 (this includes discarding the transactional footprint).

If the transaction state is Suspended, the transactional footprint is discarded (if not already discarded for a pending failure), but failure handling is deferred.

Other than the setting of CRO, execution of tabortwc. in the Non-transactional state is treated as a no-op.

\section*{Special Registers Altered:}

CRO TEXASR TFIAR TS

\section*{Transaction Abort Word Conditional Immediate \\ X-form}
tabortwci. TO,RA,SI
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & TO & RA & SI & & 846 \\
\hline 0 & 6 & & 11 & 16 & 21 \\
\hline
\end{tabular}
```

a \leftarrow EXTS((RA) 32:63)
abort }\leftarrow
CRO \leftarrow0| |SSR TS | 0
if a < EXTS(SI) \& TO
if a > EXTS(SI) \& TO
if a = EXTS(SI) \& TO2 then abort }\leftarrow
if a u< EXTS(SI) \& TO
if a >u EXTS(SI) \& TO4 then abort }\leftarrow
if abort \& (MSR
\#Transactional or Suspended
cause }\leftarrow0\times0000000
if MSR }\mp@subsup{\mp@code{TS}}{= 0b01 \& TEXASR}{FS
Discard transactional footprint
TMRecordFailure(cause)
if MSR
TMHandleFailure()

```

The tabortwci. instruction sets condition register field 0 to 0 II \(\mathrm{MSR}_{\text {TS }}\) II 0 . The contents of register \(\mathrm{RA}_{32: 63}\) are compared with the sign-extended value of the SI field. If any bit in the TO field is set to 1 and its corresponding condition is met by the result of the comparison, and the transaction state is Transactional or Suspended then the tabortwci. instruction causes transaction failure, resulting in the following:

Failure recording is performed as defined in Section 5.3 .2 , using the failure cause \(0 \times 00000001\).

If the transaction state is Transactional, failure handling is performed as defined in Section 5.3.3 (this includes discarding the transactional footprint).
If the transaction state is Suspended, the transactional footprint is discarded (if not already discarded for a pending failure), but failure handling is deferred.
Other than the setting of CRO, execution of tabortwci. in the Non-transactional state is treated as a no-op.

\section*{Special Registers Altered:}

CRO TEXASR TFIAR TS

\section*{Transaction Abort Doubleword Conditional}

X-form
tabortdc. TO,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & TO & RA & RB & & 814 & 1 \\
0 & & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

a\leftarrow(RA)
b}\leftarrow(\textrm{RB}
abort }\leftarrow
CRO \leftarrow 0 ||MSR TS || 0
if (a < b) \& TOO then abort \leftarrow
if (a>b) \& TO
if (a = b) \& TO2 then abort }\leftarrow
if (a u<b) \& TO
if (a >u b) \& TO4 then abort \leftarrow1
if abort \& (MSR RS }=0.010|MSR (TS = 0b01) the
\#Transactional or Suspended
cause \leftarrow 0x00000001
if MSR
Discard transactional footprint
TMRecordFailure(cause)
if MSR TS = 0b10 then \#Transactional
TMHandleFailure()

```

The tabortdc. instruction sets condition register field 0 to \(0\left\|\mathrm{MSR}_{\text {TS }}\right\| 0\). The contents of register RA are compared with the contents of register RB. If any bit in the TO field is set to 1 and its corresponding condition is met by the result of the comparison, and the transaction state is Transactional or Suspended, then the tabortdc. instruction causes transaction failure, resulting in the following:

Failure recording is performed as defined in Section 5.3 .2 , using the failure cause \(0 \times 00000001\).

If the transaction state is Transactional, failure handling is performed as defined in Section 5.3.3 (this includes discarding the transactional footprint).
If the transaction state is Suspended, the transactional footprint is discarded (if not already discarded for a pending failure), but failure handling is deferred.
Other than the setting of CRO, execution of tabortdc. in the Non-transactional state is treated as a no-op.

\section*{Special Registers Altered:}

CRO TEXASR TFIAR TS

Transaction Abort Doubleword Conditional Immediate

X-form
tabortdci. TO,RA, SI
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & TO & RA & SI & & 878 & 1 \\
0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

a}\leftarrow(RA
abort }\leftarrow
CRO}\leftarrow0||\mp@subsup{MSR}{TS}{}||
if a < EXTS(SI) \& TO
if a > EXTS(SI) \& TO
if a = EXTS(SI) \& TO
if a u< EXTS(SI) \& TO
if a >u EXTS(SI) \& TO
if abort \& (MSR TTS = 0b10 | MSR RS = 0b01) then
\#Transactional or Suspended
cause \leftarrow 0x00000001
if MSR
Discard transactional footprint
TMRecordFailure(cause)
if MSR TS = 0b10 then \#Transactional
TMHandleFailure()

```

The tabortdci. instruction sets condition register field 0 to \(0\left\|\mathrm{MSR}_{\mathrm{TS}}\right\| 0\). The contents of register RA are compared with the sign-extended value of the SI field. If any bit in the TO field is set to 1 and its corresponding condition is met by the result of the comparison, and the transaction state is Transactional or Suspended then the tabortdci. instruction causes transaction failure, resulting in the following:
Failure recording is performed as defined in Section 5.3 .2 , using the failure cause \(0 \times 00000001\).

If the transaction state is Transactional, failure handling is performed as defined in Section 5.3.3 (this includes discarding the transactional footprint).

If the transaction state is Suspended, the transactional footprint is discarded (if not already discarded for a pending failure), but failure handling is deferred.

Other than the setting of CRO, execution of tabortdci. in the Non-transactional state is treated as a no-op.

\section*{Special Registers Altered:}

CRO TEXASR TFIAR TS

\section*{Transaction Suspend or Resume X-form}
tsr. L
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{1}{|c|}{31} & /// & L & \multicolumn{1}{|c|}{\(/ / /\)} & \multicolumn{1}{c|}{\(/ / /\)} & & 750 & 1 \\
0 & & 6 & 10 & 11 & 16 & 21 & \\
\hline
\end{tabular}
\(\mathrm{CRO} \leftarrow 0\left\|\operatorname{MSR}_{\mathrm{TS}}\right\| 0\)

\section*{if \(L=0\) then}
if \(\begin{aligned} \mathrm{MSR}_{\mathrm{TS}} & =0 \mathrm{~b} 10 \text { then } \\ \mathrm{MSR}_{\mathrm{TS}} & \leftarrow 0 \mathrm{~b} 01\end{aligned}\)
else
if \(\mathrm{MSR}_{\mathrm{TS}}=0 \mathrm{~b} 01\) \#Suspended \(\mathrm{MSR}_{\mathrm{TS}} \leftarrow 0 \mathrm{~b} 10\)
\#Transactional

The tsr. instruction sets condition register field 0 to 0 II \(\mathrm{MSR}_{\text {TS }} I I 0\). Based on the value of the L field, two variants of tsr. are used to change the transaction state.
If \(L=0\), and the transaction state is Transactional, the transaction state is set to Suspended.
If \(L=1\), and the transaction state is Suspended, the transaction state is set to Transactional.

Other than the setting of CRO, the execution of tsr. in the Non-transactional state is treated as a no-op.

\section*{Special Registers Altered:}

\section*{CRO TS}

\section*{Programming Note}

When resuming a transaction that has encountered failure while in the Suspended state, failure handling is performed after the execution of tresume. and no later than the next failure synchronizing event.

\section*{Extended Mnemonics}

Examples of extended mnemonics for Transaction Suspend or Resume.

Transaction Check
X-form
tcheck BF
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & BF & // & I/I & I/I & & 718 \\
\hline 0 & & 6 & 9 & 11 & 16 & 21 \\
\hline
\end{tabular}
if \(M S R_{T S}=0 \mathrm{~b} 10 \mid \mathrm{MSR}_{\mathrm{TS}}=0 \mathrm{~b} 01\) then \#Transactional
\#or Suspended
for each load caused by an instruction following
the outer tbegin and preceding this tcheck
if (Load instruction was executed in \(T\) state with TEXASR \(_{\text {ROT }}=0\) or accessing a location previously stored transactionally) | (Load instruction was executed in S state with TEXASR \(_{\text {ROT }}=0\) and accessed a location previously accessed transactionally)| (Load instruction was executed in S state with TEXASR \(_{\text {ROT }}=1\) and accessed a location previously stored transactionally)
then wait until load has been performed with respect to all processors and mechanisms
CR field \(\mathrm{BF} \leftarrow\) TDOOMED || MSR \(_{\mathrm{TS}}| | 0\)

If the transaction state is Transactional or Suspended, the tcheck instruction ensures that all loads that are caused by instructions that follow the outer tbegin. instruction and precede the tcheck instruction and satisfy one of the following properties, have been performed with respect to all processors and mechanisms.
■ The load is caused by an instruction that was executed in Transactional state, either while TEXASRROT=0 or accessing a location previously stored transactionally.
- The load is caused by an instruction that was executed in Suspended state while TEXASR ROT \(^{=}=0\) and accesses a location that was accessed transactionally.
■ The load is caused by an instruction that was executed in Suspended state while TEXASR ROT \(=1\) and accesses a location that was stored transactionally.
The tcheck instruction then copies the TDOOMED bit into bit 0 of CR field BF, copies MSR Ts to bits 1:2 of CR field \(B F\), and sets bit 3 of \(C R\) field \(B F\) to 0.
Other than the setting of \(C R\) field \(B F\), execution of tcheck in the Non-transactional state is treated as a no-op.
Special Registers Altered:
CR field BF

\section*{Programming Note}

One use of the tcheck instruction in Suspended state is to determine whether preceding loads from transactionally modified locations have returned the data the transaction stored. (If the transaction has failed, some of the loads may have returned a more recent value that was stored by a conflicting store, or may have returned the pre-transaction contents of the location.). It is important to use tcheck. between any Suspended state loads that might access transactionally modified locations and subsequent computation using the Sus-pended-state-loaded data. Otherwise, corrupt data could cause problems such as wild branches or infinite loops.

Another use of tcheck in Suspended state is to determine whether the contents of storage, as seen in Suspended state, are consistent with the transaction succeeding -- e.g., whether no location that has been accessed transactionally (stored transactionally, for ROTs), and has been seen in Suspended state, has been subject to a conflict thus far. (A location is seen in Suspended state either by being loaded in Suspended state or by being loaded in Transactional state and the value (or a value derived therefrom) passed, in a register, into Suspended state.)
A use of tcheck in Transactional state is to determine whether the transaction still has the potential to succeed.
Note that tcheck provides an instantaneous check on the integrity of a subset of the accesses performed within a transaction. tcheck is not a failure synchronizing mechanism. Even if no accesses follow the tcheck, there may still be latent failures that haven't been recorded, for example caused by accesses that tcheck does not wait for, by external conflicts that will happen in the future, or simply by time of flight to the failure detection mechanism for operations that have already been performed.

\section*{Programming Note}

The tcheck instruction can return 1 in bit 0 of CR field BF before the failure has been recorded in TEXASR and TFIAR.

\section*{Programming Note}

The tcheck instruction may cause pipeline synchronization. As a result, programs that use tcheck excessively may perform poorly.

Version 3.0

\section*{Chapter 6. Time Base}

The Time Base (TB) is a 64-bit register (see Figure 9) containing a 64-bit unsigned integer that is incremented periodically as described below.
\begin{tabular}{ll|l|}
\hline & TBU & \multicolumn{1}{c|}{ TBL } \\
\hline 0 & \multicolumn{2}{|c|}{32} \\
& \\
Field & Description \\
TBU & Upper 32 bits of Time Base \\
TBL & Lower 32 bits of Time Base
\end{tabular}

Figure 9. Time Base
The Time Base monotonically increments until its value becomes 0xFFFF_FFFF__FFFF_FFFF ( \(2^{64}-1\) ); at the next increment its value becomes 0x0000_0000_0000_0000. There is no interrupt or other indication when this occurs.

The suggested frequency at which the time base increments is 512 MHz , however, variation from this rate is allowed provided the following requirements are met.
- The contents of the Time Base differ by no more than +/- four counts from what they would be if they incremented at the required frequency.
- Bit 63 of the Time Base is set to 1 between \(30 \%\) and \(70 \%\) of the time over any time interval of at least 16 counts.

The Power ISA does not specify a relationship between the frequency at which the Time Base is updated and other frequencies, such as the CPU clock or bus clock. The Time Base update frequency is not required to be constant. What is required, so that system software can keep time of day and operate interval timers, is one of the following.
- The system provides an (implementation-dependent) interrupt to software whenever the update frequency of the Time Base changes, and a means to determine what the current update frequency is.
- The update frequency of the Time Base is under the control of the system software.

\section*{Programming Note}

If the operating system initializes the Time Base on power-on to some reasonable value and the update frequency of the Time Base is constant, the Time Base can be used as a source of values that increase at a constant rate, such as for time stamps in trace entries.

Even if the update frequency is not constant, values read from the Time Base are monotonically increasing (except when the Time Base wraps from \(2^{64}-1\) to 0 ). If a trace entry is recorded each time the update frequency changes, the sequence of Time Base values can be post-processed to become actual time values.

Successive readings of the Time Base may return identical values.

\subsection*{6.1 Time Base Instructions}

\section*{Move From Time Base}

XFX-form
```

mftb RT,TBR
[Phased-Out]

```
\begin{tabular}{|l|l|ll|l|l|}
\hline 31 & \multicolumn{2}{|c|}{ RT } & \multicolumn{2}{c|}{ tbr } & \multicolumn{2}{c|}{371} & \(/\) \\
0 & & & 11 & & 21 \\
31 \\
\hline
\end{tabular}

This instruction behaves as if it were an mfspr instruction; see the mfspr instruction description in Section 3.3.17 of Book I.

\section*{Special Registers Altered:}

None

\section*{Extended Mnemonics:}

Extended mnemonics for Move From Time Base:
\begin{tabular}{|c|c|c|c|}
\hline \multicolumn{2}{|l|}{Extended:} & \multicolumn{2}{|l|}{Equivalent to:} \\
\hline mftb & Rx & mftb mfspr & \[
\begin{gathered}
\mathrm{Rx}, 268 \\
\mathrm{Rx}, 268
\end{gathered}
\] \\
\hline mftbu & Rx & mftb mfspr & \[
\begin{gathered}
\mathrm{Rx}, 269 \\
\mathrm{Rx}, 269
\end{gathered}
\] \\
\hline
\end{tabular}

\section*{Programming Note}

New programs should use mfspr instead of mftb to access the Time Base.

\section*{Programming Note}
mftb serves as both a basic and an extended mnemonic. The Assembler will recognize an mftb mnemonic with two operands as the basic form, and an mftb mnemonic with one operand as the extended form. In the extended form the TBR operand is omitted and assumed to be 268 (the value that corresponds to TB).

\section*{Programming Note}

The mfspr instruction can be used to read the Time Base on all processors that comply with Version 2.01 of the architecture or with any subsequent version.

It is believed that the mfspr instruction can be used to read the Time Base on most processors that comply with versions of the architecture that precede Version 2.01. Processors for which mfspr cannot be used to read the Time Base include the following.
```

- 601
- POWER3

```
(601 implements neither the Time Base nor mftb, but depends on software using mftb to read the Time Base, so that the attempt causes the Illegal Instruction error handler to be invoked and thereby permits the operating system to emulate the Time Base.)

\section*{Programming Note}

Since the update frequency of the Time Base is imple-mentation-dependent, the algorithm for converting the current value in the Time Base to time of day is also implementation-dependent.

As an example, assume that the Time Base increments at the constant rate of 512 MHz . (Note, however, that programs should allow for the possibility that some implementations may not increment the least-significant 4 bits of the Time Base at a constant rate.) What is wanted is the pair of 32 -bit values comprising a POSIX standard clock: \({ }^{1}\) the number of whole seconds that have passed since 00:00:00 January 1, 1970, UTC, and the remaining fraction of a second expressed as a number of nanoseconds.

Assume that:
■ The value 0 in the Time Base represents the start time of the POSIX clock (if this is not true, a simple 64-bit subtraction will make it so).
■ The integer constant ticks_per_sec contains the value \(512,000,000\), which is the number of times the Time Base is updated each second.
■ The integer constant \(n s\) _adj contains the value
\[
\frac{1,000,000,000}{512,000,000} \times 2^{32} / 2=4194304000
\]
which is the number of nanoseconds per tick of the Time Base, multiplied by \(2^{32}\) for use in mulhwu (see below), and then divided by 2 in order to fit, as an unsigned integer, into 32 bits.

When the processor is in 64-bit mode, The POSIX clock can be computed with an instruction sequence such as this:
```

mfspr Ry,268 \# Ry = Time Base
lwz Rx,ticks_per_sec
divdu Rz,Ry,Rx \# Rz = whole seconds
stw Rz,posix_sec
mulld Rz,Rz,Rx \# Rz = quotient * divisor
sub Rz,Ry,Rz \# Rz = excess ticks
lwz Rx,ns_adj
slwi Rz,Rz,1 \# Rz = 2 * excess ticks
mulhwu Rz,Rz,Rx \# mul by (ns/tick)/2 * 232
stw Rz,posix_ns\# product[0:31] = excess ns

```

\section*{Non-constant update frequency}

In a system in which the update frequency of the Time Base may change over time, it is not possible to convert an isolated Time Base value into time of day. Instead, a Time Base value has meaning only with respect to the current update frequency and the time of day that the update frequency was last changed. Each time the update frequency changes, either the system software is notified of the change via an interrupt (see Book III), or the change was instigated by the system software itself. At each such change, the system software must compute the current time of day using the old update frequency, compute a new value of ticks_per_sec for the new frequency, and save the time of day, Time Base value, and tick rate. Subsequent calls to compute Time of Day use the current Time Base Value and the saved value.

\footnotetext{
1. Described in POSIX Draft Standard P1003.4/D12, Draft Standard for Information Technology -- Portable Operating System Interface (POSIX) -Part 1: System Application Program Interface (API) - Amendment 1: Real-time Extension [C Language]. Institute of Electrical and Electronics Engineers, Inc., Feb. 1992.
}

Version 3.0

\section*{Chapter 7. Event-Based Branch Facility}

\subsection*{7.1 Event-Based Branch Overview}

The Event-Based Branch facility allows application programs to enable hardware to change the effective address of the next instruction to be executed when certain events occur to an effective address specified by the program.

The operation of the Event-Based Branch facility is summarized as follows:
- The Event-Based Branch facility is available only when the system software has made it available. See Section 9.5 of Book III for additional information.
- When the Event-Based Branch facility is available, event-based branches are caused by event-based exceptions. Event-based exceptions can be enabled to occur by setting bits in the BESCR.
- When an event-based exception occurs, the bit in the BESCR control field corresponding to the event-based exception is set to 0 and the bit in the Event Status field in the BESCR corresponding to the event-based exception is set to 1 .
- If the global enable bit in the BESCR is set to 1 when any of the bits in the status field are set to 1 (i.e., when an event-based exception exists), an event-based branch occurs.
- The event-based branch causes the following to occur.
- The global enable bit is set to 0 .
- \(\quad\) The TS field of the BESCR is set to indicate the transactional state of the processor when the event-based branch occurred; if the processor was in transactional state when the event-based branch occurred, it is put into suspended state.
- The EBBRR is set to the effective address of the instruction that would have
attempted to execute next if no event-based branch had occurred.
- Instruction fetch and execution continues at the effective address contained in the EBBHR
- The event-based branch handler performs the necessary processing in response to the event, and then executes an rfebb instruction in order to resume execution at the instruction that would have been executed next when the event-based branch occurred. The rfebb instruction also restores the processor to the transactional state indicated by BESCR Ts . See the Programming Notes in Section 7.3 for an example sequence of operations of the event-based branch handler.

Additional information about the Event-Based Branch facility is given in Section 3.4 of Book III.

\section*{Programming Note}

Since system software controls the availability of the Event-Based Branch facility (see Section 9.5 of Book III), an interface must be provided that enables applications to request access to the facility and determine when it is available.

\section*{Programming Note}

In order to initialize the Event-Based Branch facility for Performance Monitor event-based exceptions, software performs the following operations.
- Software requests control of the Event-Based Branch facility from the system software.
- Software requests the system software to initialize the Performance Monitor as desired.
- Software sets the EBBHR to the effective address of the event-based branch handler.
- Software enables Performance Monitor event-based exceptions by setting BES\(\mathrm{CR}_{\text {PME PMEO }}=10\), and also sets \(\mathrm{MMCRO}_{\text {PMAE }} \quad\) PMAO \(=10\). See Section 9.4.4 of Book III for the description of MMCRO.
- Software sets the GE bit in the BESCR to enable event-based branches.

Initializing the Event-Based Branch facility for Load Monitored and External EBB exceptions follows a similar process except that EBB exceptons for these facilities are controlled by different bits in the BESCR.

\subsection*{7.2 Event-Based Branch Registers}

\subsection*{7.2.1 Branch Event Status and Control Register}

The Branch Event Status and Control Register (BESCR) is a 64-bit register that contains control and status information about the Event-Based Branch facility.
\begin{tabular}{|l|l|l|ll|}
\hline GE & Event Control & TS & Event Status & \\
\hline 0 & 1 & & 32 & 34
\end{tabular}

Figure 10. Branch Event Status and Control Register (BESCR)
\begin{tabular}{|l|l|}
\hline GE & Event Control \\
\hline 0 & 1 \\
31 \\
\hline
\end{tabular}

Figure 11. Branch Event Status and Control Register Upper (BESCRU)

System software controls whether or not event-based branches occur regardless of the contents of the BESCR. See Section 9.4.4 of Book III and Section 6.2.12 of Book III.

The entire BESCR can be read or written using SPR 806. Individual bits of the BESCR can be set or reset using two sets of additional SPR numbers.
- When mtspr indicates SPR 800 (Branch Event Status and Control Set, or BESCRS), the bits in BESCR which correspond to "1" bits in the source register are set to 1 ; all other bits in the BESCR are unaffected. SPR 801 (BESCRSU) provides the same capability to each of the upper 32 bits of the BESCR.
- When mtspr indicates SPR 802 (Branch Event Status and Control Reset, or BESCRR), the bits in BESCR which correspond to " 1 " bits in the source register are set to 0; all other bits in the BESCR are unaffected. SPR 803 (BESCRRU) provides the same capability to each of the upper 32 bits of the BESCR.

When mfspr indicates any of the above SPR numbers, the current value of the register is returned.

\section*{Programming Note}

Event-based branch handlers typically reset event status bits upon entry, and enable event enable bits after processing an event. Execution of rfebb then re-enables the GE bit so that additional event-based branches can occur.
\(0 \quad\) Global Enable (GE)
0 Event-based branches are disabled
1 Event-based branches are enabled.
When an event-based branch occurs, GE is set to 0 and is not altered by hardware until rfebb 1 is executed or software sets GE=1 and another event-based branch occurs.

\section*{1:31 Event Control}

\section*{1:28 Reserved}

29 Load Monitored Event-Based Exception Enable (LME)
0 Load Monitored event-based exceptions are disabled.
1 If \(\mathrm{BESCR}_{G E}=1\), Load Monitored event-based exceptions are enabled until a Load Monitored event-based exception occurs, at which time:
- LME is set to 0
- LMEO is set to 1 ;

See Section 3.2.4 of Book I for information about Load Monitored event-based exceptions.

\section*{Programming Note}

When \(B_{E S C R}^{G E}=0\), the \(L M E\) bit is ignored, and Idmx behaves as Idx without affecting this bit.

\section*{30 External Event-Based Exception} Enable (EE)
0 External event-based (EBB) exceptions are disabled.
1 External EBB exceptions are enabled until an external event-based exception occurs, at which time:
- EE is set to 0
- EEO is set to 1

External event-based exceptions exist when an external EBB input from the platform is active. See the system documentation for information about the external EBB input.

\section*{31 Performance Monitor Event-Based Exception Enable (PME)}

0 Performance Monitor event-based exceptions are disabled.
1 Performance Monitor event-based exceptions are enabled until a Performance Monitor event-based exception occurs, at which time:
- PME is set to 0
- PMEO is set to 1

See Chapter 9 of Book III for information about Performance Monitor event-based exceptions and about the effects of this bit on the Performance Monitor.

\section*{32:33 Transactional State}

When an event-based branch occurs, hardware sets this field to indicate the transactional state of the processor when the event-based branch occurred.
The values and their associated meanings are as follows.
00 Non-transactional
01 Suspended
10 Transactional
11 Reserved

\section*{Programming Note}

Event-based branch handlers should not modify this field since its value is used by the processor to determine the transactional state of the processor after the \(r f e b b\) instruction is executed.

\section*{34:60Reserved}

\section*{61 Load Monitored Event-Based Excep-} tion Occurred (LMEO)
0 A Load Monitored event-based exception has not occurred since the last time software set this bit to 0 .

1 A Load Monitored event-based exception has occurred since the last time software set this bit to 0 .
This bit is set to 1 by the hardware when a Load Monitored event-based exception occurs. This bit can be set to 0 only by the mtspr instruction.

\section*{Programming Note}

Software should set this bit to 0 after handling an event-based branch due to a Load Monitored event-based exception.
62 External Event-Based Exception Occurred (EEO)
0 An external EBB exception has not occurred since the last time software set this bit to 0 .
1 An external EBB exception has occurred since the last time software set this bit to 0 .

\section*{Programming Note}

As part of processing an External EBB exception, it may also be necessary to perform additional operations to manage the external EBB input from the system. See the system documentation for details.

63 Performance Monitor Event-Based Exception Occurred (PMEO)
0 A Performance Monitor event-based exception has not occurred since the last time software set this bit to 0 .
1 A Performance Monitor event-based exception has occurred since the last time software set this bit to 0 .

This bit is set to 1 by the hardware when a Performance Monitor event-based exception occurs. This bit can be set to 0 only by the mtspr instruction.
See Chapter 9 of Book III for information about Performance Monitor event-based exceptions and about the effects of this bit on the Performance Monitor.

\section*{Programming Note}

After handling an event-based branch, software should set the "exception occurred" bit(s) corresponding to the event-based exception(s) that have occurred to 0 . See the Programming Notes in Section 7.3 for additional information.

\subsection*{7.2.2 Event-Based Branch Handler Register}

The Event-Based Branch Handler Register (EBBHR) is a 64-bit register register that contains the 62 most significant bits of the effective address of the instruction that is executed next after an event-based branch occurs. Bits \(62: 63\) must be available to be read and written by software.


Figure 12. Event-Based Branch Handler Register (EBBHR)

\section*{Programming Note}
```

The EBBHR can be used by software as a scratchpad register after entry into an event-based branch handler, provided that its contents are restored prior to executing rfebb 1. An example of such usage is as follows, where SPRG3 is used to contain a pointer to a storage area where context information may be saved.

```
// Save r1 in EBBHR
// Move SPRG3 to r1
// Store r2
// Copy original contents
// of r1 to r2
std r2,offset2(r1) // save original r1
// Store rest of state
// Process event(s)
// r1,r2
// of EBBHR in r2
// Restore EBBHR
// restore r1
// Return from handler
```

E:mtspr EBBHR, r1 // Save r1 in EBBHR

```
E:mtspr EBBHR, r1 // Save r1 in EBBHR
```

E:mtspr EBBHR, r1 // Save r1 in EBBHR
mfspr r1, SPRG3 // Move SPRG3 to r1
mfspr r1, SPRG3 // Move SPRG3 to r1
mfspr r1, SPRG3 // Move SPRG3 to r1
std r2, r1,offset1 // Store r2
std r2, r1,offset1 // Store r2
std r2, r1,offset1 // Store r2
mfspr EBBHR,r2 // Copy original contents
mfspr EBBHR,r2 // Copy original contents
mfspr EBBHR,r2 // Copy original contents
// of r1 to r2
// of r1 to r2
// of r1 to r2
std r2,offset2(r1) // save original r1
std r2,offset2(r1) // save original r1
std r2,offset2(r1) // save original r1
// Store rest of state
// Store rest of state
// Store rest of state
// Process event(s)
// Process event(s)
// Process event(s)
// Restore all state except
// Restore all state except
// Restore all state except
// r1,r2
// r1,r2
// r1,r2
r2 = \&E // Generate original value
r2 = \&E // Generate original value
r2 = \&E // Generate original value
// of EBBHR in r2
// of EBBHR in r2
// of EBBHR in r2
mtspr EBBHR,r2 // Restore EBBHR
mtspr EBBHR,r2 // Restore EBBHR
mtspr EBBHR,r2 // Restore EBBHR
ld r2 offset1(r1) // restore r2
ld r2 offset1(r1) // restore r2
ld r2 offset1(r1) // restore r2
ld r1 offset2(r1) // restore r1
ld r1 offset2(r1) // restore r1
ld r1 offset2(r1) // restore r1
rfebb 1 // Return from handler

```
    rfebb 1 // Return from handler
```

    rfebb 1 // Return from handler
    ```

\subsection*{7.2.3 Event-Based Branch Return}

\section*{Register}

The Event-Based Branch Return Register (EBBRR) is a 64 -bit register that contains the 62 most significant
He oz miver stgmicamt
bits of an instruction effective address as specified below.

When an event-based branch occurs, bits 0:61 of the EBBRR are set to the effective address of the instruction that the processor would have attempted to execute next if no event-based branch had occurred.

Bits 62:63 are reserved.


Figure 13. Event-Based Branch Return Register (EBBRR)

\subsection*{7.3 Event-Based Branch Instructions}

Return from Event-Based Branch
XL-form
rfebb S
\begin{tabular}{|l|l|l|l|l|l|ll|l|}
\hline 19 & \multicolumn{1}{|c|}{\(/ / /\)} & \multicolumn{1}{|c|}{\(/ / /\)} & \(/ / /\) & S & 146 & \(/\) \\
0 & & 6 & & 11 & 16 & 20 & 21 & \\
31 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \operatorname{BESCR}_{G E} \leftarrow S \\
& \text { MSR }_{\text {TS }} \leftarrow \operatorname{BESCR}_{\text {TS }} \\
& \text { NIA } \leftarrow \text { iea } \text { EBBRR }_{0: 61}| | 0 \mathrm{~b} 00
\end{aligned}
\]
\(B_{E S C R}^{G E}\) is set to \(S\). The processor is placed in the transactional state indicated by \(\mathrm{BESCR}_{\mathrm{TS}}\).

If there are no pending event-based exceptions, then the next instruction is fetched from the address \(\mathrm{EBBRR}_{0: 61}\) II 0 b 00 (when \(\mathrm{MSR}_{\mathrm{SF}}=1\) ) or \({ }^{32} 0\) II \(\operatorname{EBBRR}_{32: 61}\) II \(0 \mathrm{bO0}\) (when MSR \({ }_{\text {SF }}=0\) ). If one or more pending event-based exceptions exist, an event-based branch is generated; in this case the value placed into EBBRR by the Event-Based Branch facility is the address of the instruction that would have been executed next had the event-based branch not occurred.

See Section 3.4 of Book III for additional information about this instruction.

\section*{Special Registers Altered:}

BESCR
MSR (See Book III)

\section*{Extended Mnemonics:}

\section*{Extended:}
rfebb

Equivalent to:
rfebb 1

\section*{Programming Note}
rfebb serves as both a basic and an extended mnemonic. The Assembler will recognize an rfebb mnemonic with one operand as the basic form, and an rfebb mnemonic with no operand as the extended form. In the extended form, the S operand is omitted and assumed to be 1 .

\section*{Programming Note}

If the \(\mathrm{BESCR}_{\text {TS }}\) has been modified by software after an event-based branch occurs, an illegal transaction state transition may occur. See Chapter 3.2.2 of Book III.

\section*{Programming Note}

When an event-based branch occurs, the event-based branch handler can execute the following sequence of operations. This sequence of operations assumes that the handler routine has access to a stack or other area in memory in which state information from the main program can be stored. Note also that in this example, the handler entry point is labeled " \(E\)," \(r\) 1 and \(r 2\) are used as scratch registers, and both external EBB and Performance Monitor EBB exceptions are enabled.
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{E:Save state // This is the entry pt} \\
\hline mfspr r1, BESCR & // Check event status \\
\hline if \(r 1_{63}=1\), then & \\
\hline \multicolumn{2}{|l|}{Process PM exception} \\
\hline \multicolumn{2}{|l|}{\(r 2 \leftarrow 0 x 0000000000000001\)} \\
\hline mtspr BESCRR, r2 & //Reset PMEO status bit \\
\hline \multicolumn{2}{|l|}{\(r 2 \leftarrow 0 x 0000000100000000\)} \\
\hline \multicolumn{2}{|l|}{mtspr BESCRS, r1 //Re-enable PM exceptions} \\
\hline \multicolumn{2}{|l|}{//Note: The PMAE bit of MMCRO must also} \\
\hline // be & enabled. See Book III. \\
\hline \multicolumn{2}{|l|}{if \(\mathrm{r} 1_{62}=1\), then} \\
\hline \multicolumn{2}{|l|}{Process external exception} \\
\hline \multicolumn{2}{|l|}{\(r 2 \leftarrow 0 x 0000000000000002\)} \\
\hline \multicolumn{2}{|l|}{mtspr BESCRR, r2 //Reset EEO status bit} \\
\hline \multicolumn{2}{|l|}{\(r 2 \leftarrow 0 \times 0000000200000000\)} \\
\hline \multicolumn{2}{|r|}{// De-activate external EBB input from platform} \\
\hline \multirow[t]{6}{*}{mtspr BESCRS, r1} & //Re-enable external EBB exceptions \\
\hline & // . . \\
\hline & //Other exceptions such as \\
\hline & //Load Monitor exceptions \\
\hline & //are processed similarly. \\
\hline & // . . \\
\hline \multicolumn{2}{|l|}{Restore state} \\
\hline rfebb 1 & // return \& global enable \\
\hline
\end{tabular}

Note that before resetting the BESCR EEO , the external EBB input from the platform should be deactivated, and additional operations to manage the external EBB input may be required. See the system documentation for details.

In the above sequence, if other exceptions occur after they are enabled, another event-based branch will occur immediately after rebb is executed.

\title{
Appendix A. Assembler Extended Mnemonics
}

In order to make assembler language programs simpler to write and easier to understand, a set of extended mnemonics and symbols is provided for certain instructions. This appendix defines extended mnemonics and
symbols related to instructions defined in Book II. Assemblers should provide the extended mnemonics and symbols listed here, and may provide others.

\section*{A. 1 Data Cache Block Touch [for Store] Mnemonics}

The TH field in the Data Cache Block Touch and Data Cache Block Touch for Store instructions control the actions performed by the instructions. Extended mnemonics are provided that represent the TH value in the mnemonic rather than requiring it to be coded as a numeric operand.
\begin{tabular}{|c|c|}
\hline dcbtct RA,RB, TH & \begin{tabular}{l}
(equivalent to: dcbt for TH values of Ob00000-0b00111); \\
other TH values are invalid
\end{tabular} \\
\hline dcbtds RA, RB, TH & (equivalent to: dcbt for TH values of 0b00000 or Ob01000 - Ob01111); other TH values are invalid. \\
\hline dcbtt RA,RB & (equivalent to: dcbt for TH value of 0b10000) \\
\hline dcbna RA,RB & (equivalent to: dcbt for TH value of Ob10001) \\
\hline dcbtstct RA,RB,TH & (equivalent to: dcbtst for TH values of 0b00000 or Ob00000-0b00111); other TH values are invalid. \\
\hline dcbtstds RA,RB, TH & (equivalent to: dcbtst for TH values of 0b00000 or Ob01000-0b01111); other TH values are invalid. \\
\hline dcbtstt RA,RB & (equivalent to: dcbtst for TH value of Ob10000) \\
\hline
\end{tabular}

\section*{A. 2 Data Cache Block Flush Mnemonics}

The L field in the Data Cache Block Flush instruction controls the scope of the flush function performed by the instruction. Extended mnemonics are provided that
represent the \(L\) value in the mnemonic rather than requiring it to be coded as a numeric operand.
Note: dcbf serves as both a basic and an extended mnemonic. The Assembler will recognize a dcbf mnemonic with three operands as the basic form, and a dcbf mnemonic with two operands as the extended form. In the extended form the L operand is omitted and assumed to be 0 .
\[
\begin{array}{ll}
\text { dcbf RA,RB } & \text { (equivalent to: dcbf RA,RB,0) } \\
\text { dcbfl RA,RB } & \text { (equivalent to: dcbf RA,RB,1) } \\
\text { dcbflp RA,RB } & \text { (equivalent to: dcbf RA,RB,3) }
\end{array}
\]

\section*{A. 3 Or Mnemonics}

The three register fields in the or instruction can be used to specify a hint indicating how the processor should handle stores caused by previous Store or dcbz instructions. An extended mnemonic is supported that represents the operand values in the mnemonic rather than requiring them to be coded as numeric operands.

> miso
(equivalent to: or \(26,26,26\) )

\section*{A. 4 Load and Reserve Mnemonics}

The EH field in the Load and Reserve instructions provides a hint regarding the type of algorithm implemented by the instruction sequence being executed. Extended mnemonics are provided that allow the EH value to be omitted and assumed to be 0b0.

Note: Ibarx, Iharx, Iwarx, Idarx, and Iqarx serve as both basic and extended mnemonics. The Assembler will recognize these mnemonics with four operands as the basic form, and these mnemonics with three oper-
ands as the extended form. In the extended form the EH operand is omitted and assumed to be 0 .
\begin{tabular}{lll} 
Ibarx \(R T, R A, R B\) & (equivalent to: Ibarx & \(R T, R A, R B, 0\) ) \\
Iharx \(R T, R A, R B\) & (equivalent to: Iharx & \(R T, R A, R B, 0\) ) \\
Iwarx \(R T, R A, R B\) & (equivalent to: Iwarx & \(R T, R A, R B, 0\) ) \\
Idarx \(R T, R A, R B\) & (equivalent to: Idarx & \(R T, R A, R B, 0\) ) \\
Iqarx \(R T, R A, R B\) & (equivalent to: Iqarx & \(R T, R A, R B, 0\) )
\end{tabular}

\section*{A. 5 Synchronize Mnemonics}

The L field in the Synchronize instruction controls the scope of the synchronization function performed by the instruction. Extended mnemonics are provided that represent the \(L\) value in the mnemonic rather than requiring it to be coded as a numeric operand. Two extended mnemonics are provided for the \(L=0\) value in order to support Assemblers that do not recognize the sync mnemonic.
Note: sync serves as both a basic and an extended mnemonic. Assemblers will recognize a sync mnemonic with one operand as the basic form, and a sync mnemonic with no operand as the extended form. In the extended form the \(L\) operand is omitted and assumed to be 0 .
\begin{tabular}{llll} 
sync & (equivalent to: & sync & 0 ) \\
Iwsync & (equivalent to: & sync & \(1)\) \\
ptesync & (equivalent to: & sync & \(2)\)
\end{tabular}

\section*{A. 6 Wait Mnemonics}

The WC field in the wait instruction is reserved for future use. It may be be used in the future to indicate the condition that causes instruction execution to resume. An extended mnemonic is provided that represent the WC value in the mnemonic rather than requiring it to be coded as a numeric operand.

Note: wait serves as both a basic and an extended mnemonic. The Assembler will recognize a wait mnemonic with one operand as the basic form, and a wait mnemonic with no operands as the extended form. In the extended form the WC operand is omitted and assumed to be 0 .
wait
(equivalent to: wait 0)

\section*{A. 7 Transactional Memory Instruction Mnemics}

The A field in the Transaction End instruction controls whether the instruction ends only the current (possibly nested) transaction or the entire set of nested transactions. Extended mnemonics are provided that represent
the \(A\) value in the mnemonic rather than requiring it to be coded as a numeric operand..
\[
\begin{array}{ll}
\text { tend. } & \text { (equivalent to: tend. 0) } \\
\text { tendall. } & \text { (equivalent to: tend. 1) }
\end{array}
\]

The L field in the Transaction Suspend or Resume instruction determines how to change the transaction state. Extended mnemonics are provided that represent the \(L\) value in the mnemonic rather than requiring it to be coded as a numeric operand.
\begin{tabular}{ll} 
tsuspend. & (equivalent to: tsr. 0) \\
tresume. & (equivalent to: tsr. 1)
\end{tabular}

\section*{A. 8 Move To/From Time Base Mnemonics}

The tbr field in the Move From Time Base instruction specifies whether the instruction reads the entire Time Base or only the high-order half of the Time Base.
\begin{tabular}{rr} 
mftb \(R x\) & (equivalent to: \(m f t b R x, 268)\) \\
or: \(m f s p r, 268\) \\
mftbu \(R x\) & (equivalent to: \(m f t b R x, 269)\) \\
& or: \(m f s p r, 269\)
\end{tabular}

\section*{A. 9 Return From Event-Based Branch Mnemonic}

The S field in the Return from Event-Based Branch instruction specifies the value to which the instruction sets the GE field in the BESCR. Extended mnemonics are provided that represent the \(S\) value in the mnemonic rather than requiring it to be coded as a numeric operand.
rfebb (equivalent to: rfebb 1)
Note: rfebb serves as both a basic and an extended mnemonic. The Assembler will recognize this mnemonic with one operand as the basic form, and this mnemonic with no operands as the extended form. In the extended form the \(S\) operand is omitted and assumed to be 1.

\title{
Appendix B. Programming Examples for Sharing Storage
}

\begin{abstract}
This appendix gives examples of how dependencies and the Synchronization instructions can be used to control storage access ordering when storage is shared between programs.

Many of the examples use extended mnemonics (e.g., bne, bne-, cmpw) that are defined in Appendix \(C\) of Book I.

Many of the examples use the Load And Reserve and Store Conditional instructions, in a sequence that begins with a Load And Reserve instruction and ends with a Store Conditional instruction (specifying the same storage location as the Load Conditional) followed by a Branch Conditional instruction that tests whether the Store Conditional instruction succeeded.
\end{abstract}

In these examples it is assumed that contention for the shared resource is low; the conditional branches are optimized for this case by using " + " and "-" suffixes appropriately.
The examples deal with words; they can be used for doublewords by changing all word-specific mnemonics to the corresponding doubleword-specific mnemonics (e.g., Iwarx to Idarx, cmpw to cmpd).

In this appendix it is assumed that all shared storage locations are in storage that is Memory Coherence Required, and that the storage locations specified by Load And Reserve and Store Conditional instructions are in storage that is neither Write Through Required nor Caching Inhibited.

\section*{B. 1 Atomic Update Primitives}

This section gives examples of how the Load And Reserve and Store Conditional instructions can be used to emulate atomic read/modify/write operations.

An atomic read/modify/write operation reads a storage location and writes its next value, which may be a function of its current value, all as a single atomic operation. The examples shown provide the effect of an atomic read/modify/write operation, but use several instructions rather than a single atomic instruction.

\section*{Fetch and No-op}

The "Fetch and No-op" primitive atomically loads the current value in a word in storage.
In this example it is assumed that the address of the word to be loaded is in GPR 3 and the data loaded are returned in GPR 4.
loop:
lwarx r4,0,r3 \#load and reserve
stwcx. r4,0,r3 \#store old value if
\# still reserved
bne- loop \#loop if lost reservation
Note:
1. The stwcx., if it succeeds, stores to the target location the same value that was loaded by the preceding Iwarx. While the store is redundant with respect to the value in the location, its success ensures that the value loaded by the Iwarx is still the current value at the time the stwcx. is executed.

\section*{Fetch and Store}

The "Fetch and Store" primitive atomically loads and replaces a word in storage.
In this example it is assumed that the address of the word to be loaded and replaced is in GPR 3, the new value is in GPR 4, and the old value is returned in GPR 5.
```

loop:
lwarx r5,0,r3 \#load and reserve
stwcx. r4,0,r3 \#store new value if
\# still reserved
bne- loop loop if lost reservation

```

\section*{Fetch and Add}

The "Fetch and Add" primitive atomically increments a word in storage.

In this example it is assumed that the address of the word to be incremented is in GPR 3, the increment is in GPR 4, and the old value is returned in GPR 5.
```

loop:
lwarx r5,0,r3 \#load and reserve
add r0,r4,r5\#increment word
stwcx. r0,0,r3 \#store new value if still res'ved
bne- loop \#loop if lost reservation

```

\section*{Fetch and AND}

The "Fetch and AND" primitive atomically ANDs a value into a word in storage.

In this example it is assumed that the address of the word to be ANDed is in GPR 3, the value to AND into it is in GPR 4, and the old value is returned in GPR 5.
```

loop:
lwarx r5,0,r3 \#load and reserve
and r0,r4,r5\#AND word
stwcx. r0,0,r3 \#store new value if still res'ved
bne- loop \#loop if lost reservation

```

Note:
1. The sequence given above can be changed to perform another Boolean operation atomically on a word in storage, simply by changing the and instruction to the desired Boolean instruction (or, xor, etc.).

\section*{Test and Set}

This version of the "Test and Set" primitive atomically loads a word from storage, sets the word in storage to a nonzero value if the value loaded is zero, and sets the EQ bit of CR Field 0 to indicate whether the value loaded is zero.

In this example it is assumed that the address of the word to be tested is in GPR 3, the new value (nonzero) is in GPR 4, and the old value is returned in GPR 5.
```

loop:
lwarx r5,0,r3 \#load and reserve
cmpwi r5,0 \#done if word not equal to 0
bne- exit
stwcx. r4,0,r3 \#try to store non-0
bne- loop \#loop if lost reservation
exit: ...

```

\section*{Compare and Swap}

The "Compare and Swap" primitive atomically compares a value in a register with a word in storage, if they are equal stores the value from a second register into the word in storage, if they are unequal loads the word from storage into the first register, and sets the EQ bit of CR Field 0 to indicate the result of the comparison.

In this example it is assumed that the address of the word to be tested is in GPR 3, the comparand is in GPR 4 and the old value is returned there, and the new value is in GPR 5.
loop:
lwarx r6,0,r3 \#load and reserve
cmpw r4,r6 \#1st 2 operands equal?
bne- exit \#skip if not
stwcx. r5, 0, r3 \#store new value if still res'ved
bne- loop \#loop if lost reservation
exit:
mr r4,r6 \#return value from storage
Notes:
1. The semantics given for "Compare and Swap" above are based on those of the IBM System/370 Compare and Swap instruction. Other architectures may define a Compare and Swap instruction differently.
2. "Compare and Swap" is shown primarily for pedagogical reasons. It is useful on machines that lack the better synchronization facilities provided by Iwarx and stwcx.. A major weakness of a Sys-tem/370-style Compare and Swap instruction is that, although the instruction itself is atomic, it checks only that the old and current values of the word being tested are equal, with the result that programs that use such a Compare and Swap to control a shared resource can err if the word has been modified and the old value subsequently restored. The sequence shown above has the same weakness.
3. In some applications the second bne- instruction and/or the \(\boldsymbol{m r}\) instruction can be omitted. The bne- is needed only if the application requires that if the EQ bit of CR Field 0 on exit indicates "not equal" then (r4) and (r6) are in fact not equal. The \(\boldsymbol{m r}\) is needed only if the application requires that if the comparands are not equal then the word from storage is loaded into the register with which it was compared (rather than into a third register). If either or both of these instructions is omitted, the resulting Compare and Swap does not obey System/370 semantics.

\section*{B. 2 Lock Acquisition and Release, and Related Techniques}

This section gives examples of how dependencies and the Synchronization instructions can be used to imple-
ment locks, import and export barriers, and similar constructs.

\section*{B.2.1 Lock Acquisition and Import Barriers}

An "import barrier" is an instruction or sequence of instructions that prevents storage accesses caused by instructions following the barrier from being performed before storage accesses that acquire a lock have been performed. An import barrier can be used to ensure that a shared data structure protected by a lock is not accessed until the lock has been acquired. A sync instruction can be used as an import barrier, but the approaches shown below will generally yield better performance because they order only the relevant storage accesses.

\section*{B.2.1.1 Acquire Lock and Import Shared Storage}

If Iwarx and stwcx. instructions are used to obtain the lock, an import barrier can be constructed by placing an isync instruction immediately following the loop containing the Iwarx and stwcx.. The following example uses the "Compare and Swap" primitive to acquire the lock.

In this example it is assumed that the address of the lock is in GPR 3, the value indicating that the lock is free is in GPR 4, the value to which the lock should be set is in GPR 5, the old value of the lock is returned in GPR 6, and the address of the shared data structure is in GPR 9.
```

loop:
lwarx r6,0,r3,1 \#load lock and reserve
cmpw r4,r6 \#skip ahead if
bne- wait \# lock not free
stwcx. r5,0,r3 \#try to set lock
bne- loop \#loop if lost reservation
isync \#import barrier
lwz r7,data1(r9)\#load shared data
wait... \#wait for lock to free

```

The hint provided with Iwarx indicates that after the program acquires the lock variable (i.e., stwcx. is successful), it will release it (i.e., store to it) prior to another program attempting to modify it.
The second bne- does not complete until CRO has been set by the stwcx. The stwcx. does not set CR0 until it has completed (successfully or unsuccessfully). The lock is acquired when the stwcx. completes successfully. Together, the second bne- and the subse-
quent isync create an import barrier that prevents the load from "data1" from being performed until the branch has been resolved not to be taken.
If the shared data structure is in storage that is neither Write Through Required nor Caching Inhibited, an Iwsync instruction can be used instead of the isync instruction. If Iwsync is used, the load from "data1" may be performed before the stwcx.. But if the stwcx. fails, the second branch is taken and the Iwarx is re-executed. If the stwcx. succeeds, the value returned by the load from "data1" is valid even if the load is performed before the stwcx., because the Iwsync ensures that the load is performed after the instance of the Iwarx that created the reservation used by the successful stwcx..

\section*{B.2.1.2 Obtain Pointer and Import Shared Storage}

If Iwarx and stwcx. instructions are used to obtain a pointer into a shared data structure, an import barrier is not needed if all the accesses to the shared data structure depend on the value obtained for the pointer. The following example uses the "Fetch and Add" primitive to obtain and increment the pointer.

In this example it is assumed that the address of the pointer is in GPR 3, the value to be added to the pointer is in GPR 4, and the old value of the pointer is returned in GPR 5.
```

loop:
lwarx r5,0,r3 \#load pointer and reserve
add r0,r4,r5\#increment the pointer
stwcx. r0,0,r3 \#try to store new value
bne- loop \#loop if lost reservation
lwz r7,data1(r5) \#load shared data

```

The load from "data1" cannot be performed until the pointer value has been loaded into GPR 5 by the Iwarx. The load from "data1" may be performed before the \(\boldsymbol{s t w c x}\). But if the stwcx. fails, the branch is taken and the value returned by the load from "data1" is discarded. If the stwcx. succeeds, the value returned by the load from "data1" is valid even if the load is performed before the stwcx., because the load uses the pointer value returned by the instance of the Iwarx that created the reservation used by the successful stwcx..

An isync instruction could be placed between the bneand the subsequent Iwz, but no isync is needed if all accesses to the shared data structure depend on the value returned by the Iwarx.

\section*{B.2.2 Lock Release and Export Barriers}

An "export barrier" is an instruction or sequence of instructions that prevents the store that releases a lock from being performed before stores caused by instructions preceding the barrier have been performed. An export barrier can be used to ensure that all stores to a shared data structure protected by a lock will be performed with respect to any other processor before the store that releases the lock is performed with respect to that processor.

\section*{B.2.2.1 Export Shared Storage and Release Lock}

A sync instruction can be used as an export barrier independent of the storage control attributes (e.g., presence or absence of the Caching Inhibited attribute) of the storage containing the shared data structure. Because the lock must be in storage that is neither Write Through Required nor Caching Inhibited, if the shared data structure is in storage that is Write Through Required or Caching Inhibited a sync instruction must be used as the export barrier.

In this example it is assumed that the shared data structure is in storage that is Caching Inhibited, the address of the lock is in GPR 3, the value indicating that the lock is free is in GPR 4, and the address of the shared data structure is in GPR 9.
```

stw r7,data1(r9)\#store shared data (last)
sync \#export barrier
stw r4,lock(r3)\#release lock

```

The sync ensures that the store that releases the lock will not be performed with respect to any other processor until all stores caused by instructions preceding the sync have been performed with respect to that processor.

\section*{B.2.2.2 Export Shared Storage and Release Lock using Iwsync}

If the shared data structure is in storage that is neither Write Through Required nor Caching Inhibited, an Iwsync instruction can be used as the export barrier. Using Iwsync rather than sync will yield better performance in most systems.

In this example it is assumed that the shared data structure is in storage that is neither Write Through Required nor Caching Inhibited, the address of the lock is in GPR 3, the value indicating that the lock is free is in GPR 4, and the address of the shared data structure is in GPR 9.
```

stw r7,data1(r9)\#store shared data (last)
lwsync \#export barrier
stw r4,lock(r3)\#release lock

```

The Iwsync ensures that the store that releases the lock will not be performed with respect to any other processor until all stores caused by instructions preceding the Iwsync have been performed with respect to that processor.

\section*{B.2.3 Safe Fetch}

If a load must be performed before a subsequent store (e.g., the store that releases a lock protecting a shared data structure), a technique similar to the following can be used.

In this example it is assumed that the address of the storage operand to be loaded is in GPR 3, the contents of the storage operand are returned in GPR 4, and the address of the storage operand to be stored is in GPR 5.
```

1wz r4,0(r3)\#load shared data
cmpw r4,r4 \#set CR0 to "equal"
bne- \$-8 \#branch never taken
stw r7,0(r5)\#store other shared data

```

An alternative is to use a technique similar to that described in Section B.2.1.2, by causing the stw to depend on the value returned by the Iwz and omitting the cmpw and bne-. The dependency could be created by ANDing the value returned by the Iwz with zero and then adding the result to the value to be stored by the \(\boldsymbol{s t w}\). If both storage operands are in storage that is neither Write Through Required nor Caching Inhibited, another alternative is to replace the cmpw and bnewith an Iwsync instruction.

\section*{B. 3 List Insertion}

This section shows how the Iwarx and stwcx. instructions can be used to implement simple insertion into a singly linked list. (Complicated list insertion, in which multiple values must be changed atomically, or in which the correct order of insertion depends on the contents of the elements, cannot be implemented in the manner shown below and requires a more complicated strategy such as using locks.)

The "next element pointer" from the list element after which the new element is to be inserted, here called the "parent element", is stored into the new element, so that the new element points to the next element in the list; this store is performed unconditionally. Then the address of the new element is conditionally stored into the parent element, thereby adding the new element to the list.

In this example it is assumed that the address of the parent element is in GPR 3, the address of the new element is in GPR 4, and the next element pointer is at offset 0 from the start of the element. It is also assumed that the next element pointer of each list element is in a reservation granule separate from that of the next element pointer of all other list elements.
loop:
lwarx r2,0,r3 \#get next pointer
stw r2,0(r4)\#store in new element
lwsync or sync \#order stw before stwcx
stwcx. r4,0,r3 \#add new element to list
bne- loop \#loop if stwcx. failed
In the preceding example, if two list elements have next element pointers in the same reservation granule then, in a multiprocessor, "livelock" can occur. (Livelock is a state in which processors interact in a way such that no processor makes forward progress.)
If it is not possible to allocate list elements such that each element's next element pointer is in a different reservation granule, then livelock can be avoided by using the following, more complicated, sequence.
```

    lwz r2,0(r3)#get next pointer
    loop1:
mr r5,r2 \#keep a copy
stw r2,0(r4)\#store in new element
sync \#order stw before stwcx.
and before lwarx
loop2:
1warx r2,0,r3 \#get it again
cmpw r2,r5 \#loop if changed (someone
bne- loop1 \# else progressed)
stwcx. r4,0,r3 \#add new element to list
bne- loop2 \#loop if failed

```

In the preceding example, livelock is avoided by the fact that each processor re-executes the stw only if some other processor has made forward progress.

\section*{B. 4 Notes}

The following notes apply to Section B. 1 through Section B. 3 .
1. To increase the likelihood that forward progress is made, it is important that looping on Iwarx/stwex. pairs be minimized. For example, in the "Test and Set" sequence shown in Section B.1, this is achieved by testing the old value before attempting the store; were the order reversed, more stwcx. instructions might be executed, and reservations might more often be lost between the Iwarx and the stwcx.
2. The manner in which Iwarx and stwcx. are communicated to other processors and mechanisms, and between levels of the storage hierarchy within a given processor, is implementation-dependent. In some implementations performance may be improved by minimizing looping on a Iwarx instruction that fails to return a desired value. For example, in the "Test and Set" sequence shown in Section B.1, if the programmer wishes to stay in the loop until the word loaded is zero, he could change the "bne- exit" to "bne- loop". However, in some implementations better performance may be obtained by using an ordinary Load instruction to do the initial checking of the value, as follows.
```

loop:
lwz r5,0(r3)\#load the word
cmpwi r5,0 \#loop back if word
bne- loop \# not equal to 0
lwarx r5,0,r3 \#try again, reserving
cmpwi r5,0 \# (likely to succeed)
bne- loop
stwcx.r4,0,r3 \#try to store non-0
bne- loop \#loop if lost reserv'n

```
3. In a multiprocessor, livelock is possible if there is a Store instruction (or any other instruction that can clear another processor's reservation; see Section 1.7.4.1) between the Iwarx and the stwcx. of a Iwarx/stwcx. loop and any byte of the storage location specified by the Store is in the reservation granule. For example, the first code sequence shown in Section B. 3 can cause livelock if two list elements have next element pointers in the same reservation granule.

\section*{B. 5 Transactional Lock Elision}

This section illustrates the use of the Transactional Memory facility to implement transactional lock elision (TLE), in which lock-based critical sections are speculatively executed as a transaction without first acquiring a lock. This locking protocol is an alternative to the routines described above, yielding increased concurrency when the lock that guards a critical section is frequently unnecessary.

\section*{B.5.1 Enter Critical Section}

The following example shows the entry point to a critical section using transactional lock elision. The entry code starts a transaction using the tbegin. instruction and checks whether the transaction was aborted or not. If not, it checks whether the lock is free or not. If the lock is found to be free, the thread proceeds to execute the critical section.

In this example it is assumed that the address of the lock is in GPR 3, and the value indicating that the lock is free is in GPR 4. The handling of cases of transaction abort and busy lock are described in subsequent examples.
```

tle_entry:
tbegin. \#Start TLE transaction
beq- tle_abort \#Handle TLE transaction abort
lwz r6,0(r3)
cmpw r6,r4
bne- busy_lock \#If not, handle lock busy case
critical_section1:

```

\section*{B.5.2 Handling Busy Lock}

In the event that the lock is already held, by either another thread or the current thread, the transaction is aborted using the tabort instruction, using a soft-ware-defined code TLE_BUSY_LOCK indicating the cause of the abort. The abort returns control to the beq following tbegin. in the critical section entrance sequence, allowing for an abort handler to react appropriately.
```

busy_lock:
li r3, TLE_BUSY_LOCK
tabort r3
\#Abort TLE transaction

```

\section*{B.5.3 Handling TLE Abort}

A TLE transaction may fail for one of a variety of causes, persistent and transient. Persistent causes are certain-or at least highly likely-to cause future attempts to execute the same transaction to fail. However, for transient causes, it is possible that the failure cause may not be re-encountered in a subsequent attempt. Thus, persistent aborts are handled by taking a non-transactional path that involves the actual acquisition of the lock, while transient aborts retry the critical section using TLE.

The following example illustrates the handling of aborts in TLE. It is assumed that the address of the lock is in

GPR 3. The immediate value of the andis. instruction selects the Failure Persistent bit in the upper half of TEXASR to be tested.
```

tle_abort:
mfspr r4, TEXASRU \# Read high-order half
\# of TEXASR
andis. r5,r4,0x0100 \# determine whether failure
\# is likely to be persistent
bne tle_acquire_lock \#Persistent, acquire lock
\#enter critical sec
b tle_entry \#Transient, try TLE again

```

This example can be extended to keep track of the number of transient aborts and fall back on the acquisition of the lock after the number of transient failures reaches some threshold. It can also be extended to handle reentrant locks. Acquisition of TLE locks is described in a subsequent example.

\section*{B.5.4 TLE Exit Section Critical Path}

The following example illustrates the instruction sequence used to exit a TLE critical section. The CRO value set by tend. indicates whether the current thread was in a transaction. If so, the exited critical section was entered speculatively, and the transaction is ended. If not, the execution takes a path to release the lock.

Release of an acquired TLE lock is described in a subsequent example.
```

tle_exit:
tend. \#End the current trans-
\#action, if any
bng- tle_release_lock \#Release lock, if was
\#not in a transaction

```

\section*{B.5.5 Acquisition and Release of TLE Locks}

The steps for acquiring and releasing a lock associated with a TLE critical section are identical to those for acquiring and releasing conventional locks that are not elided, as described in Section B.2.1.1 and Section B.2.2 respectively.

\section*{Programming Note}

A future version of the architecture will revise the isync and Iwsync instruction descriptions to make them consistent with the use of these instructions, as shown in Section B.2.1.1, to acquire a lock associated with a TLE critical section.

Version 3.0

\section*{Book III:}

Power ISA Operating Environment Architecture

\title{
Chapter 1. Introduction
}

\subsection*{1.1 Overview}

Chapter 1 of Book I describes computation modes, document conventions, a general systems overview, instruction formats, and storage addressing. This chapter augments that description as necessary for the Power ISA Operating Environment Architecture.

\subsection*{1.2 Document Conventions}

The notation and terminology used in Book I apply to this Book also, with the following substitutions.

■ For "system alignment error handler" substitute "Alignment interrupt".

■ For "system data storage error handler" substitute "Data Storage interrupt", "Hypervisor Data Storage interrupt", or "Data Segment interrupt", as appropriate.

■ For "system error handler" substitute "interrupt".
■ For "system floating-point enabled exception error handler" substitute "Floating-Point Enabled Exception type Program interrupt".
- For "system illegal instruction error handler" substitute "Hypervisor Emulation Assistance interrupt".
■ For "system instruction storage error handler" substitute "Instruction Storage interrupt", "Hypervisor Instruction Storage interrupt", or "Instruction Segment interrupt", as appropriate.
- For "system privileged instruction error handler" substitute "Privileged Instruction type Program interrupt".
- For "system service program" substitute "System Call interrupt" or "System Call Vectored interrupt", as appropriate.
■ For "system trap handler" substitute "Trap type Program interrupt".

■ For "system facility unavailable error handler" substitute "Facility Unavailable interrupt" or "Hypervisor Facility Unavailable interrupt."

\subsection*{1.2.1 Definitions and Notation}

The definitions and notation given in Book I and Book II are augmented by the following.

■ Threaded processor, single-threaded processor, thread
A threaded processor implements one or more "threads", where a thread corresponds to the Book I/II concept of "processor". That is, the definition of "thread" is the same as the Book I definition of "processor", and "processor" as used in Books I and II can be thought of as either a single-threaded processor or as one thread of a multi-threaded processor. Except where the meaning is clear in context or the number of threads does not matter, the only unqualified uses of "processor" in Book III are in resource names (e.g. Processor Identification Register); such uses should be regarded as meaning "threaded processor". The threads of a multi-threaded processor typically share certain resources, such as the hardware components that execute certain kinds of instructions (e.g., Fixed-Point instructions), certain caches, the address translation mechanism, and certain hypervisor resources.

■ real page
A unit of real storage that is aligned at a boundary that is a multiple of its size. The real page size is 4 KB .

■ context of a program
The state (e.g., privilege and relocation) in which the program executes. The context is controlled by the contents of certain System Registers, such as the MSR and PTCR, of certain lookaside buffers, such as the SLB and TLB, and of the Page Table.
- performed

The definition of "performed" given in Section 1.1 of Book II is extended to apply to implicit storage accesses and to invalidations of entries in caches of information derived from address translation tables, as follows.
- The definition of "load is performed" applies to accsses for performing address translation.
- The definition of "store is performed" applies to accesses for recording reference and change information.
- A TLB entry invalidation by thread T1 is performed with respect to thread T2 when the instruction that requested the invalidation has caused the specified entry, if present, to be made invalid in T2's TLB, and similarly for invalidations of entries in other caches of information derived from tables used in address translation.

\section*{- exception}

An error, unusual condition, or external signal, that may set a status bit and may or may not cause an interrupt, depending upon whether the corresponding interrupt is enabled.

\section*{- interrupt}

The act of changing the machine state in response to an exception, as described in Chapter 6. "Interrupts" on page 1047.

\section*{- trap interrupt}

An interrupt that results from execution of a Trap instruction.
- Additional exceptions to the rule that the thread obeys the sequential execution model, beyond those described in Section 2.2 of Book I and in the bullet defining "program order" in Section 1.1 of Book II, are the following.
- A System Reset or Machine Check interrupt may occur. The determination of whether an instruction is required by the sequential execution model is not affected by the potential occurrence of a System Reset or Machine Check interrupt. (The determination is affected by the potential occurrence of any other kind of interrupt.)
- A context-altering instruction is executed (Chapter 11. "Synchronization Requirements for Context Alterations" on page 1127). The context alteration need not take effect until the required subsequent synchronizing operation has occurred.
- A Reference and Change bit is updated by the thread. The update need not be performed with respect to that thread until the required subsequent synchronizing operation has occurred.
- A Branch instruction is executed and the branch is taken. The update of the Come-From Address Register (see Section 8.2 of Book III) need not occur until a subsequent context synchronizing operation has occurred.
- An mtgsr is executed and an interrupt or event-based branch occurs before the mtspr
sequence following mtgsr has finished executing. The contents of SPRs that are the targets of mtspr instructions between the point of interruption or EBB and the end of the mtspr sequence may be altered.
■ "must"
If hypervisor software violates a rule that is stated using the word "must" (e.g., "this field must be set to 0 "), and the rule pertains to the contents of a hypervisor resource, to executing an instruction that can be executed only in hypervisor state, or to accessing storage in real addressing mode, the results are undefined, and may include altering resources belonging to other partitions, causing the system to "hang", etc.
- hardware

Any combination of hard-wired implementation, emulation assist, or interrupt for software assistance. In the last case, the interrupt may be to an architected location or to an implementa-tion-dependent location. Any use of emulation assists or interrupts to implement the architecture is implementation-dependent.

\section*{- hypervisor privileged}

A term used to describe an instruction or facility that is available only when the thread is in hypervisor state.
- privileged state and supervisor mode

Used interchangeably to refer to a state in which privileged facilities are available.
- problem state and user mode

Used interchangeably to refer to a state in which privileged facilities are not available.
■ /, I/, I/I, ... denotes a field that is reserved in an instruction, in a register, or in an architected storage table.

■ ?, ??, ???, ... denotes a field that is implementa-tion-dependent in an instruction, in a register, or in an architected storage table.

\subsection*{1.2.2 Reserved Fields}

Book l's description of the handling of reserved bits in System Registers, and of reserved values of defined fields of System Registers, applies also to the SLB. Book I's description of the handling of reserved values of defined fields of System Registers applies also to architected storage tables (e.g., the Page Table).
Some fields of certain architected storage tables may be written to automatically by the hardware, e.g., Reference and Change bits in the Page Table. When the hardware writes to such a table, the following rules are obeyed.
- Unless otherwise stated, no defined field other than the one(s) specifically being updated are modified.
- Contents of reserved fields are either preserved or written as zero.

\section*{Programming Note}

Software should set reserved fields in the SLB and in architected storage tables to zero, because these fields may be assigned a meaning in some future version of the architecture.

\subsection*{1.3 General Systems Overview}

The hardware contains the sequencing and processing controls for instruction fetch, instruction execution, and interrupt action. Most implementations also contain data and instruction caches. Instructions that the processing unit can execute fall into the following classes:

■ instructions executed in the Branch Facility
- instructions executed in the Fixed-Point Facility
- instructions executed in the Floating-Point Facility
- instructions executed in the Vector Facility

Almost all instructions executed in the Branch Facility, Fixed-Point Facility, Floating-Point Facility, and Vector Facility are nonprivileged and are described in Book I. Book II may describe additional nonprivileged instructions (e.g., Book II describes some nonprivileged instructions for cache management). Instructions related to the privileged state, control of hardware resources, control of the storage hierarchy, and all other privileged instructions are described here or are implementation-dependent.

\subsection*{1.4 Exceptions}

The following augments the exceptions defined in Book I that can be caused directly by the execution of an instruction:
■ the execution of a floating-point instruction when \(M S R_{F P}=0\) (Floating-Point Unavailable interrupt)
- an attempt to modify a hypervisor resource when the thread is in privileged but non-hypervisor state (see Chapter 2), or an attempt to execute a hyper-visor-only instruction (e.g., tlbie) when the thread is in privileged but non-hypervisor state
- the execution of a traced instruction (Trace interrupt)
- the execution of a Vector instruction when the vector facility is unavailable (Vector Unavailable interrupt)

\subsection*{1.5 Synchronization}

The synchronization described in this section refers to the state of the thread that is performing the synchronization.

\subsection*{1.5.1 Context Synchronization}

An instruction or event is context synchronizing if it satisfies the requirements listed below. Such instructions and events are collectively called context synchronizing operations. The context synchronizing operations are the isync instruction, the System Linkage instructions, the \(\boldsymbol{m t m s r}[d]\) instructions with \(L=0\), and most interrupts (see Section 6.4).
1. The operation causes instruction dispatching (the issuance of instructions by the instruction fetching mechanism to any instruction execution mechanism) to be halted.
2. The operation is not initiated or, in the case of isync, does not complete, until all instructions that precede the operation have completed to a point at which they have reported all exceptions they will cause.
3. The operation ensures that the instructions that precede the operation will complete execution in the context (privilege, relocation, storage protection, etc.) in which they were initiated, except that the operation has no effect on the context in which the associated Reference and Change bit updates are performed.
4. If the operation directly causes an interrupt (e.g., \(\boldsymbol{s c}\) directly causes a System Call interrupt) or is an interrupt, the operation is not initiated until no exception exists having higher priority than the exception associated with the interrupt (see Section 6.9).
5. The operation ensures that the instructions that follow the operation will be fetched and executed in the context established by the operation. (This requirement dictates that any prefetched instructions be discarded and that any effects and side effects of executing them out-of-order also be discarded, except as described in Section 5.5, "Performing Operations Out-of-Order".)

\section*{Programming Note}

A context synchronizing operation is necessarily execution synchronizing; see Section 1.5.2.

Unlike the Synchronize instruction, a context synchronizing operation does not affect the order in which storage accesses are performed.

Item 2 permits a choice only for isync (and sync and ptesync; see Section 1.5.2) because all other execution synchronizing operations also alter context.

\subsection*{1.5.2 Execution Synchronization}

An instruction is execution synchronizing if it satisfies items 2 and 3 of the definition of context synchronization (see Section 1.5.1). sync and ptesync are treated like isync with respect to item 2. The execution synchronizing instructions are sync, ptesync, the \(\boldsymbol{m t m s r}[d]\) instructions with \(L=1\), and all context synchronizing instructions.
\[
\begin{aligned}
& \text { Programming Note } \\
& \text { Unlike a context synchronizing operation, an execu- } \\
& \text { tion synchronizing instruction does not ensure that } \\
& \text { the instructions following that instruction will exe- } \\
& \text { cute in the context established by that instruction. } \\
& \text { This new context becomes effective sometime after } \\
& \text { the execution synchronizing instruction completes } \\
& \text { and before or at a subsequent context synchroniz- } \\
& \text { ing operation. }
\end{aligned}
\]

\title{
Chapter 2. Logical Partitioning (LPAR) and Thread Control
}

\subsection*{2.1 Overview}

The Logical Partitioning (LPAR) facility permits threads and portions of real storage to be assigned to logical collections called partitions, such that a program executing on a thread in one partition cannot interfere with any program executing on a thread in a different partition. This isolation can be provided for both problem state and privileged non-hypervisor state programs, by using a layer of trusted software, called a hypervisor program (or simply a "hypervisor"), and the resources provided by this facility to manage system resources. (A hypervisor is a program that runs in hypervisor state; see below.)

The number of partitions supported is implementa-tion-dependent.

A thread is assigned to one partition at any given time. A thread can be assigned to any given partition without consideration of the physical configuration of the system (e.g., shared registers, caches, organization of the storage hierarchy), except that threads that share certain hypervisor resources may need to be assigned to the same partition; see Section 2.6. The registers and facilities used to control Logical Partitioning are listed below and described in the following subsections.

Except in the following subsections, references to the "operating system" in this document include the hypervisor unless otherwise stated or obvious from context.

\subsection*{2.2 Logical Partitioning Control Register (LPCR)}

The contents of the LPCR control a number of aspects of the operation of the thread with respect to a logical partition. Below are shown the bit definitions for the LPCR.
\begin{tabular}{ll|} 
Bit & Description \\
0:3 & Virtualization Control (VC) \\
& Controls the virtualization of partition memory. \\
& This field contains three subfields, VPM, ISL, \\
and KBV. Accesses that are initiated in hyper- \\
& visor state (i.e., MSR MV PR=0b10) are per- \\
formed as if VC=Ob0000.
\end{tabular}

0:1 Virtualized Partition Memory (VPM)
This field controls whether VPM mode is enabled as specified below. (See Section 5.7.3.3 and Section 5.7.2, "Virtualized Partition Memory (VPM) Mode" for additional information on VPM mode.)

\section*{Bit Description}

I 0 Reserved

1 This bit controls whether VPM mode is enabled when address translation is enabled
0 - VPM mode disabled
1 - VPM mode enabled

\section*{Programming Note}

VPM1 must be set to zero by hypervisors that want to receive interrupts from applications running directly under them as DSIs (instead of HDSIs). See Section 6.5.3, "Data Storage Interrupt" for more information.

2 Ignore SLB Large Page Specification (ISL)
Controls whether ISL mode is enabled as specified below.

0 - ISL mode disabled
1 - ISL mode enabled
When ISL mode is enabled and address translation is enabled, address translation is performed as if the contents of SLB \(_{\text {LIILP }}\) and PRTE STPS were 0b000. When address
translation is disabled, the setting of the ISL bit has no effect. ISL mode has no effect on SLB, TLB, and ERAT entry invalidations caused by slbie, slbieg, slbia, tlbie, and tlbiel.

\section*{Programming Note}

Specifying that LIILP=0b000 in PATE PS when VPM mode is enabled has the same effect on address translation when translation is disabled as enabling ISL mode when translation is enabled.

ISL mode is needed when translation is enabled because translation uses the SLB, and the contents of the SLB are accessible to the operating system and should not be modified by the hypervisor. ISL mode is not needed when translation is disabled since Virtual Real Mode address translation uses PATE \(_{\text {PS }}\), which is not visible to the operating system and is in complete control of the hypervisor.

\section*{3 Key-Based Virtualization (KBV)}

Controls whether Key-Based Virtualization is enabled as specified below.
\(0-\) KBV is disabled
\(1-\mathrm{KBV}\) is enabled
When KBV is enabled, Virtual Page Class Key Storage Protection exceptions that occur on operand accesses when \(\mathrm{VPM}_{1}=0\) cause Hypervisor Data Storage interrupts.

\section*{Programming Note}

Key-Based Virtualization provides an efficient means for the hypervisor to intercept storage references, e.g. MMIO, that must be emulated. (The corresponding behavior for instruction fetching is not desired.) Virtual Page Class Key Storage Protection exceptions not handled by the hypervisor should be reflected to the operating system at its Data Storage interrupt vector with the hypervisor having set DSISR \(_{42}\).

\section*{4:8 Reserved}

\section*{Default Prefetch Depth (DPFD)}

The DPFD field is used as the default prefetch depth for data stream prefetching when \(D_{S C R}^{\text {DPFD }}=0\); see page 844.
12:16 Reserved

17:19

17 Hypervisor Virtualization Exit Enable
0 When the stop instruction is executed with PSSCR \(_{\mathrm{EC}}=1\), Hypervisor Virtualization exceptions are not enabled to cause exit from power-saving mode.
1 When the stop instruction is executed with PSSCR \(_{\text {EC }}=1\), Hypervisor Virtualization exceptions are enabled to cause exit from power-saving mode.
18:19 Reserved
20:37 Reserved
38 Interrupt Little-Endian (ILE)
The contents of the ILE bit are copied into \(\mathrm{MSR}_{\mathrm{LE}}\) by interrupts that set MSR \({ }_{\mathrm{HV}}\) to 0 (see Section 6.5), to establish the Endian mode for the interrupt handler.
39:40 Alternate Interrupt Location (AIL)
Controls the effective address offset, or alternate effective address for System Call Vectored, of the interrupt handler and the relocation mode in which it begins execution for all interrupts except those subject to the overrides described below.
0 The interrupt is taken with \(\mathrm{MSR}_{\mathrm{IR} \text { DR }}=\) Ob00 and no effective address offset or alternate effective address.
1 Reserved
2 The interrupt is taken with \(M S R_{I R} D R=\) Ob11. If the interrupt is not System Call Vectored, an effective address offset of 0x0000_0000_0001_8000 is applied. System Call Vectored uses an alternate effective address of 0x0000_0000_0001_7 || LEV || 0b0_0000.
3 The interrupt is taken with \(M S R_{I R} D R=\) Ob11. If the interrupt is not System Cal Vectored, an effective address offset of 0xc000_0000_0000_4000 is applied. System Call Vectored uses an alternate effective address of 0xc000_0000_0000_3 || LEV || 0b0_0000.
Machine Check, System Reset, and Hypervisor Maintenance interrupts are taken as if \(\mathrm{LPCR}_{\text {AIL }}=0\). In the remainder of this definition, "other interrupts" means interrupts other than these three.

Other interrupts that occur when \(\mathrm{MSR}_{\mathrm{IR}}=0\) or \(\mathrm{MSR}_{\mathrm{DR}}=0\), are taken as if \(\mathrm{LPCR}_{\text {AIL }}=0\).

When the software receiving other interrupts uses Radix Tree translation and the interrupts occur when \(M S R_{I_{R}}=1\) and \(M S R_{D R}=1\), the interrupts are taken as if \(\mathrm{LPCR}_{\text {AIL }}=3\). This includes interrupts taken by the hypervisor
when PATE \(\mathrm{ERR}=1\) and by the operating system (or a nested hypervisor) when PATE \(\mathrm{GR}^{=1}\).
When the hypervisor receiving the other interrupts uses HPT translation and the interrupts have caused a transition from \(\mathrm{MSR}_{\mathrm{HV}}=0\) to \(\mathrm{MSR}_{\mathrm{HV}}=1\), the interrupts are taken as if \(\operatorname{LPCR}_{\text {AIL }}=0\).

\section*{- Programming Note}

One of the purposes of the AIL field is to provide relocation for interrupts that occur while an application is running with \(\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR}^{=}=0 \mathrm{~b} 11\) under a "bare metal" operating system (i.e., an operating system that runs in hypervisor state), such as KVM.

Oniine (ONL)

0 The PURR and SPURR do not increment.
1 The PURR and SPURR increment.

\section*{Programming Note}

Typically, the hypervisor sets the ONL bit to 0 when the thread is not in a power saving mode, is not performing useful work, and is available for use. The hypervisor may take the state of the ONL bit into account when making course-grain load balancing and power management decisions.
\(46 \quad\) Large Decrementer (LD)
0 Large Decrementer mode is not enabled.
1 Large Decrementer mode is enabled.
See Section 7.4 for additional information.
47:51 Power-saving mode Exit Cause Enable (Lower Section) (PECE \(L\) )
Privileged Doorbell Exit Enable
0 When the stop instruction is executed with PSSCR \(_{E C}=1\), Directed Privileged Doorbell exceptions are not enabled to cause exit from power-saving mode
1 When the stop instruction is executed with PSSCR \(_{E C}=1\), Directed Privileged Doorbell exceptions are enabled to cause exit from power-saving mode.

\section*{Hypervisor Doorbell Exit Enable}

0 When the stop instruction is executed with PSSCR \(_{\text {EC }}=1\), Directed Hypervisor Doorbell exceptions are not enabled to cause exit from power-saving mode
1 When the stop instruction is executed with PSSCR \(_{\text {EC }}=1\), Directed Hypervisor Doorbell exceptions are enabled to cause exit from power-saving mode.

\section*{External Exit Enable}

0 When the stop instruction is executed with PSSCR \(_{\text {EC }}=1\), External exceptions are not enabled to cause exit from power-saving mode.
1 When the stop instruction is executed with PSSCR \(_{\text {EC }}=1\), External exceptions are enabled to cause exit from power-saving mode.

\section*{Decrementer Exit Enable}

0 When the stop instruction is executed with PSSCR \(_{\text {EC }}=1\), Decrementer exceptions are not enabled to cause exit from power-saving mode.
1 When the stop instruction is executed with PSSCR \(_{E C}=1\), Decrementer exceptions are enabled to cause exit from power-saving mode. (Decrementer exceptions do not occur if the state of the Decrementer is not maintained and updated as
if the thread was not in power-saving mode.)

\section*{51 Other Exit Enable}

0 When the stop instruction is executed with PSSCR \({ }_{E C}=1\), Machine Check, Hypervisor Maintenance, and certain implemen-tation-specific exceptions are not enabled to cause exit from power-saving mode.
1 When the stop instruction is executed with PSSCR \(_{E C}=1\), Machine Check, Hypervisor Maintenance, and certain implemen-tation-specific exceptions are enabled to cause exit from power-saving mode.

If the state of the PECE field is lost during power-saving mode, implementations must provide the means to exit power-saving mode upon the occurrence of a System Reset exception and any of the exceptions that were enabled by the PECE field when the stop instruction was executed. In addition, they may also exit power-saving mode on exceptions that were disabled by the PECE field as well. See Section 6.5.1 and Section 6.5.2 for additional information about exit from power-saving mode.

\section*{52 Mediated External Exception Request (MER)}

0 A Mediated External exception is not requested.
1 A Mediated External exception is requested.
The exception effects of this bit are said to be consistent with the contents of this bit if one of the following statements is true.
- \(\quad \operatorname{LPCR}_{\text {MER }}=1\) and a Mediated External exception exists.
- \(\quad \operatorname{LPCR}_{\text {MER }}=0\) and a Mediated External exception does not exist.
A context synchronizing instruction or event that is executed or occurs when LPCR \(_{\text {MER }}=0\) ensures that the exception effects of LPCR \(_{\text {MER }}\) are consistent with the contents of LPCR \(_{\text {MER }}\). Otherwise, when an instruction changes the contents of LPCR MER , the exception effects of LPCR \(_{\text {MER }}\) become consistent with the new contents of LPCR \(_{\text {MER }}\) reasonably soon after the change.

\section*{Programming Note}

LPCR \(_{\text {MER }}\) provides a means for the hypervisor to direct an external exception to a partition independent of the partition's \(\mathrm{MSR}_{\mathrm{EE}}\) setting. (When \(M S R_{E E}=0\), it is inappropriate for the hypervisor to deliver the exception.) Using LPCR \({ }_{\text {MER }}\), the partition can be interrupted upon enabling external interrupts. Without using LPCR \(_{\text {MER }}\), the hypervisor must check the state of \(\mathrm{MSR}_{\text {EE }}\) whenever it gets control, which will result in less timely delivery of the exception to the partition.

Guest Translation Shootdown Enable (GTSE)

Controls whether the operating system is permitted to use tlbie and slbieg directly, or must issue a system call to the hypervisor.
0 Guest is not permitted to use tlbie, slbieg, tlbsync, and slbsync.
1 Guest is permitted to use tlbie, slbieg, tlbsync, and slbsync.

\section*{Translation Control (TC)}

0 The secondary Page Table search is enabled.
1 The secondary Page Table search is disabled.

Reserved
Hypervisor External Interrupt Control (HEIC)

0 Direct External interrupts can occur in Hypervisor state.
1 Direct External interrupts cannot occur in hypervisor state.

\section*{Programming Note}

By setting HEIC=1, the Hypervisor Interrupt Virtualization handler can prevent External interrupts from occurring during the Hypervisor Virtualization interrupt handler. See Section 6.5.7.1.

60 Logical Partitioning Environment Selector (LPES)

0 External interrupts set the HSRRs, set \(M_{S R}\) to 1 , and leave \(M S R_{R I}\) unchanged.
1 External interrupts set the SRRs, set \(\mathrm{MSR}_{\mathrm{RI}}\) to 0 , and leave \(\mathrm{MSR}_{\mathrm{HV}}\) unchanged.

> - Programming Note
> LPES \(=1\) should be used by operating systems not running under a hypervisor, so that external interrupts are directed to the SRRs rather than to the HSRRs.

\section*{Programming Note}

In versions of the architecture that precede Version 2.07, LPES was a two-bit field, in which the second bit controlled significant aspects of storage accessing and interrupt handling.

\section*{61 Reserved}

62 Hypervisor Virtualization Interrupt Conditionally Enable (HVICE)
0 Hypervisor Virtualization interrupts are disabled.
1 Hypervisor Virtualization interrupts are enabled if permitted by \(\mathrm{MSR}_{\text {EE }}, \mathrm{MSR}_{\mathrm{HV}}\), and MSR \({ }_{\text {PR }}\); see Section 6.5.21.

Hypervisor Decrementer Interrupt Conditionally Enable (HDICE)

0 Hypervisor Decrementer interrupts are disabled.
1 Hypervisor Decrementer interrupts are enabled if permitted by \(\mathrm{MSR}_{\mathrm{EE}}, \mathrm{MSR}_{\mathrm{HV}}\), and MSR \(_{\text {PR }}\); see Section 6.5.12 on page 1073.

See Section 6.5 on page 1060 for a description of how the setting of LPES affects the processing of interrupts.

\subsection*{2.3 Hypervisor Real Mode Offset Register (HRMOR)}

The layout of the Hypervisor Real Mode Offset Register (HRMOR) is shown in Figure 1 below.
\begin{tabular}{|l|ll|}
\hline\(/ /\) & & \multicolumn{1}{c|}{ HRMO } \\
\hline \(0 \quad 4\) & & \\
Bits & Name & Description \\
\(4: 63\) & HRMO & Real Mode Offset
\end{tabular}

Figure 1. Hypervisor Real Mode Offset Register All other fields are reserved.

The supported HRMO values are the non-negative multiples of \(2^{r}\), where \(r\) is an implementation-dependent value and \(12 \leq r \leq 26\).

The contents of the HRMOR affect how some storage accesses are performed as described in Section 5.7.3 on page 984 and Section 5.7.5 on page 987.

\subsection*{2.4 Logical Partition Identification Register (LPIDR)}

The layout of the Logical Partition Identification Register (LPIDR) is shown in Figure 2 below.


Figure 2. Logical Partition Identification Register
The contents of the LPIDR identify the partition to which the thread is assigned, affecting some aspects of translation and interrupt delivery. The number of LPIDR bits supported is implementation-dependent.

\section*{Programming Note}

Radix tree translation assigns special meaning to LPID=0, specifically indicating the hypervisor's own partition. When HR=1, LPIDR should not be set to zero except when \(\mathrm{MSR}_{\mathrm{HV}}=1\).

HPT translation provides special functionality for LPID=0 when \(H V=1\), as described in Section 5.7.12.4, to support the execution of a "bare metal" operating system (an operating system that runs in hypervisor state). Because the partition with LPID=0 has a somewhat different programming model from the other partitions, LPID=0 may be undesirable for use by a guest operating system running under an HPT hypervisor.

\subsection*{2.5 Processor Compatibility Register (PCR)}

The layout of the Processor Compatibility Register (PCR) is shown in Figure 3 below.


Figure 3. Processor Compatibility Register
I
Each defined bit in the PCR controls whether certain instructions, SPRs, and other related facilities are available in problem state. Except as specified elsewhere in this section, the PCR has no effect on facilities when the thread is not in problem state. Facilities that are made unavailable by the PCR are treated as follows when the thread is in problem state.
- Instructions are treated as illegal instructions,
- SPRs are treated as if they were not defined for the implementation,
- The "reserved SPRs" (see Section 1.3.3 of Book I) are treated as not defined for the implementation,
- Fields in instructions are treated as if they were 0s,
- bits in system registers read back 0s, and mtspr operations have no effect on their values.

\section*{Programming Note}

When a bit in a system register is made unavailable by the PCR, mtspr operations performed on the register in problem state have no effect on the value of the bit regardless of the privilege state in which the register may subsequently be read.

\section*{I}

A PCR bit may also determine how an instruction field value is interpreted or may define other behavior as specified in the bit definitions below.

The PCR has no effect on the setting of the MSR and [H]SRR1 by interrupts (and of the Count Register by the System Call Vectored interrupt), and by the rfscv, [h]rfid and mtmsr[d] instructions, except as specified elsewhere in this section.

\section*{Programming Note}

Because the PCR does not prevent mtspr, rfscv, [h]rfid, and mtmsr[d] instructions from setting bits in system registers that the PCR will make unavailable after a transition to problem state, these instructions may cause interrupts in a variety of unexpected ways. For example, consider an operating system that sets SRR1 such that rfid returns to problem state with MSR[TS] nonzero. A TM Bad Thing interrupt will result, despite that TM is made unavailable by the PCR.

Similarly, the PCR does not prevent rfebb instructions from setting bits in system registers that the PCR has made unavailable in problem state, and thus changes to \(B E S C R_{\text {Ts }}\) made by privileged code have the potential to subsequently cause illegal transaction state transitions when rfebb is executed in problem state, resulting in the occurrence of TM Bad Thing type Program interrupts.

When facilities that have enable bits in the MSR, FSCR, HFSCR, or MMCRO are made unavailable by the value in the PCR, they become unavailable in problem state as specified above regardless of whether they are enabled by the corresponding MSR, FSCR, HFSCR, or

MMCRO bit; facility availability interrupts (e.g. [Hypervisor] Facility Available, Vector Unavailable, etc.) do not occur as a result of problem state accesses even if the corresponding field in the MSR, [H]FSCR, or MMCRO makes them unavailable in problem state.

\section*{Programming Note}

Facilities that can be disabled in problem state by the PCR that also have enable bits in either the MSR or [H]FSCR include Transactional Memory, the BHRB instructions, event-based branch instructions, TAR, DSCR at SPR 3, SIER, MMCR2, the event-based branch instructions, and certain Float-ing-Point, Vector, and VSX instructions. When any of these facilities are made unavailable in problem state by the PCR, the corresponding [Hypervisor] Facility Unavailable, Floating-Point Unavailable, Vector, or VSX unavailable interrupts do not occur when the facility is accessed in problem state. Note, however, that the PCR does not affect privileged accesses, and thus any Hypervisor Facility Unavailable, Floating-Point Unavailable, Vector unavailable, or VSX unavailable interrupts that are specified to occur as a result of privileged accesses occur regardless of the PCR value.

The bit definitions for the PCR are shown below.
Bit Description
0:59 Reserved
\(60 \quad\) Version 2.07 (v2.07)
This bit controls the availability, in problem state, of the following instructions, facilities, and behaviors that were newly available in problem state in the version of the architecture subsequent to Version 2.07.
- The instructions listed in Table 1
- The implicit setting of \(X_{E R}{ }_{\text {OV32 }}\) and XER \(_{\text {CA32 }}\) by Fixed-Point Arithmetic instructions
- LMRR, LMSER
- scv

0 The instructions, behaviors, and facilities listed above are available in problem state.

1 The instructions, behaviors, and facilities listed above are unavailable in problem state.
\begin{tabular}{|c|c|}
\hline Mnemonic & Instruction Name \\
\hline addpcis & Add PC Immediate Shifted Prefix \\
\hline bcdcfn. & Decimal Convert From National \\
\hline bcdcfsq. & Decimal Convert From Signed Qword \\
\hline bcdcfz. & Decimal Convert From Zoned \\
\hline bcdcpsgn & Decimal CopySign \\
\hline bcdctn. & Decimal Convert To National \\
\hline bcdctsq. & Decimal Convert To Signed Qword \\
\hline bcdctz. & Decimal Convert To Zoned \\
\hline bcds. & Decimal Shift \\
\hline bcdsetsgn. & Decimal Set Sign \\
\hline bcdsr. & Decimal Shift and Round \\
\hline bcdtrunc. & Decimal Truncate \\
\hline bcdus. & Decimal Unsigned Shift \\
\hline bcdutrunc. & Decimal Unsigned Truncate \\
\hline cmpeqb & Compare Equal Byte \\
\hline cmprb & Compare Ranged Byte \\
\hline cnttzd[.] & Count Trailing Zeros Dword \\
\hline cnttzw[.] & Count Trailing Zeros Word \\
\hline copy & Copy \\
\hline cp_abort & CP_Abort \\
\hline darn & Deliver a Random Number \\
\hline dtstsfi & DFP Test Significance Immediate \\
\hline dtstsfiq & DFP Test Significance Immediate Quad \\
\hline extswsli[.] & Extend Sign Word and Shift Left Immediate \\
\hline Idat & Load Doubleword Atomic \\
\hline Idmx & Load Monitored Indexed \\
\hline Iwat & Load Word Atomic \\
\hline Ixsd & Load VSX Scalar Dword \\
\hline Ixsibzx & Load VSX Scalar as Integer Byte \& Zero Indexed \\
\hline Ixsihzx & Load VSX Scalar as Integer Hword \& Zero Indexed \\
\hline Ixssp & Load VSX Scalar Single \\
\hline Ixv & Load VSX Vector \\
\hline lxvb16x & Load VSX Vector Byte*16 Indexed \\
\hline Ixvh8x & Load VSX Vector Halfword*8 Indexed \\
\hline Ixv| & Load VSX Vector with Length \\
\hline Ixvl| & Load VSX Vector Left-justified with Length \\
\hline Ixvwsx & Load VSX Vector Word \& Splat Indexed \\
\hline Ixvx & Load VSX Vector Indexed \\
\hline maddhd & Multiply-Add High Dword \\
\hline maddhdu & Multiply-Add High Dword Unsigned \\
\hline maddld & Multiply-Add Low Dword \\
\hline mcrxrx & Move XER to CR Extended \\
\hline mfvsrld & Move From VSR Lower Dword \\
\hline modsd & Modulo Signed Dword \\
\hline modsw & Modulo Signed Word \\
\hline
\end{tabular}

Version 3.0
\begin{tabular}{|c|c|}
\hline Mnemonic & Instruction Name \\
\hline modud & Modulo Unsigned Dword \\
\hline moduw & Modulo Unsigned Word \\
\hline mtvsrdd & Move To VSR Double Dword \\
\hline mtvsrws & Move To VSR Word \& Splat \\
\hline paste[.] & Paste \\
\hline paste[.] & Paste \\
\hline setb & Set Boolean \\
\hline stdat & Store Doubleword Atomic \\
\hline stwat & Store Word Atomic \\
\hline stxsd & Store VSX Scalar Dword \\
\hline stxsibx & Store VSX Scalar as Integer Byte Indexed \\
\hline stxsihx & Store VSX Scalar as Integer Hword Indexed \\
\hline stxssp & Store VSX Scalar Single \\
\hline stxv & Store VSX Vector \\
\hline stxvb16x & Store VSX Vector Byte*16 Indexed \\
\hline stxvh8x & Store VSX Vector Halfword*8 Indexed \\
\hline stxvl & Store VSX Vector with Length \\
\hline stxvll & Store VSX Vector Left-justified with Length \\
\hline stxvx & Store VSX Vector Indexed \\
\hline sync 3 & Sync (L=3) Copy-Paste Sync \\
\hline vabsdub & Vector Absolute Difference Unsigned Byte \\
\hline vabsduh & Vector Absolute Difference Unsigned Hword \\
\hline vabsduw & Vector Absolute Difference Unsigned Word \\
\hline vbpermd & Vector Bit Permute Dword \\
\hline vclzisbb & Vector Count Leading Zero Least-Significant Bits Byte \\
\hline vcmpneb[.] & Vector Compare Not Equal Byte \\
\hline vcmpneh[.] & Vector Compare Not Equal Hword \\
\hline vcmpnew[.] & Vector Compare Not Equal Word \\
\hline vcmpnezb[.] & Vector Compare Not Equal or Zero Byte \\
\hline vcmpnezh[.] & Vector Compare Not Equal or Zero Hword \\
\hline vcmpnezw[.] & Vector Compare Not Equal or Zero Word \\
\hline vctzb & Vector Count Trailing Zeros Byte \\
\hline vctzd & Vector Count Trailing Zeros Dword \\
\hline vctzh & Vector Count Trailing Zeros Hword \\
\hline vctzlsbb & Vector Count Trailing Zero Least-Significant Bits Byte \\
\hline vctzw & Vector Count Trailing Zeros Word \\
\hline vextractd & Vector Extract Dword \\
\hline vextractub & Vector Extract Unsigned Byte \\
\hline vextractuh & Vector Extract Unsigned Hword \\
\hline vextractuw & Vector Extract Unsigned Word \\
\hline vextsb2d & Vector Extend Sign Byte To Dword \\
\hline vextsb2w & Vector Extend Sign Byte To Word \\
\hline vextsh2d & Vector Extend Sign Hword To Dword \\
\hline vextsh2w & Vector Extend Sign Hword To Word \\
\hline vextsw2d & Vector Extend Sign Word To Dword \\
\hline vextublx & Vector Extract Unsigned Byte Left-Indexed \\
\hline vextubrx & Vector Extract Unsigned Byte Right-Indexed \\
\hline vextuhlx & Vector Extract Unsigned Hword Left-Indexed \\
\hline
\end{tabular}
\begin{tabular}{|c|c|}
\hline Mnemonic & Instruction Name \\
\hline vextuhrx & Vector Extract Unsigned Hword Right-Indexed \\
\hline vextuwlx & Vector Extract Unsigned Word Left-Indexed \\
\hline vextuwrx & Vector Extract Unsigned Word Right-Indexed \\
\hline vinsertb & Vector Insert Byte \\
\hline vinsertd & Vector Insert Dword \\
\hline vinserth & Vector Insert Hword \\
\hline vinsertw & Vector Insert Word \\
\hline vmul10cuq & Vector Multiply-by-10 \& write Carry Unsigned Qword \\
\hline vmul10ecuq & Vector Multiply-by-10 Extended \& write Carry Unsigned Qword \\
\hline vmul10euq & Vector Multiply-by-10 Extended Unsigned Qword \\
\hline vmul10uq & Vector Multiply-by-10 Unsigned Qword \\
\hline vnegd & Vector Negate Dword \\
\hline vnegw & Vector Negate Word \\
\hline vpermr & Vector Permute Right-indexed \\
\hline vprtybd & Vector Parity Byte Dword \\
\hline vprtybq & Vector Parity Byte Qword \\
\hline vprtybw & Vector Parity Byte Word \\
\hline vrldmi & Vector Rotate Left Dword then Mask Insert \\
\hline vrldnm & Vector Rotate Left Dword then AND with Mask \\
\hline vrlwmi & Vector Rotate Left Word then Mask Insert \\
\hline vrlwnm & Vector Rotate Left Word then AND with Mask \\
\hline vslv & Vector Shift Left Variable \\
\hline vsrv & Vector Shift Right Variable \\
\hline wait & Wait \\
\hline xsabsqp & VSX Scalar Quad-Precision Absolute \\
\hline xsaddqp[o] & VSX Scalar Quad-Precision Add [\& round to Odd] \\
\hline xscmpexpdp & VSX Scalar Double-Precision Compare Exponents \\
\hline xscmpexpqp & VSX Scalar Quad-Precision Compare Exponents \\
\hline xscmpoqp & VSX Scalar Quad-Precision Compare Ordered \\
\hline xscmpuqp & VSX Scalar Quad-Precision Compare Unordered \\
\hline xscpsgnqp & VSX Scalar Quad-Precision CopySign \\
\hline xscvdpqp & VSX Scalar Quad-Precision Convert From Double-Precision \\
\hline xscvhpsp & VSX Scalar Convert Half-Precision to Double-Precision \\
\hline xscvqpdp[0] & VSX Scalar round \& Convert Quad-Precision to Double-Precision [using round to Odd] \\
\hline xscvqpsdz & VSX Scalar truncate \& Convert Quad-Precision to Signed Dword \\
\hline xscvqpswz & VSX Scalar truncate \& Convert Quad-Precision to Signed Word \\
\hline xscvqpudz & VSX Scalar truncate \& Convert Quad-Precision to Unsigned Dword \\
\hline xscvqpuwz & VSX Scalar truncate \& Convert Quad-Precision to Unsigned Word \\
\hline xscvsdqp & VSX Scalar Convert Signed Dword format to Quad-Precision format \\
\hline xscvsphp & VSX Scalar round \& Convert Double-Precision to Half-Precision \\
\hline xscvudqp & VSX Scalar Convert Unsigned Dword format to Quad-Precision format \\
\hline xsdivqp[o] & VSX Scalar Quad-Precision Divide [\& round to Odd] \\
\hline xsiexpdp & VSX Scalar Double-Precision Insert Exponent \\
\hline xsiexpqp & VSX Scalar Quad-Precision Insert Exponent \\
\hline xsmaddqp[0] & VSX Scalar Quad-Precision Multiply-Add [\& round to Odd] \\
\hline xsmsubqp[o] & VSX Scalar Quad-Precision Multiply-Subtract [\& round to Odd] \\
\hline xsmulqp[o] & VSX Scalar Quad-Precision Multiply [\& round to Odd] \\
\hline
\end{tabular}

\section*{Table 1: Instructions Controlled by the V 2.07 Bit}

\section*{Version 3.0}
\begin{tabular}{|l|l|}
\hline Mnemonic & Instruction Name \\
\hline xsnabsqp & VSX Scalar Quad-Precision Negative Absolute \\
\hline xsnegqp & VSX Scalar Quad-Precision Negate \\
\hline xsnmaddqp[o] & VSX Scalar Quad-Precision Negative Multiply-Add [\& round to Odd] \\
\hline xsnmsubqp[o] & VSX Scalar Quad-Precision Negative Multiply-Subtract [\& round to Odd] \\
\hline xsrqpi & VSX Scalar Round to Quad-Precision Integer \\
\hline xsrqpxp & VSX Scalar Quad-Precision Round to Double-Extended-Precision \\
\hline xssqrtqp[o] & VSX Scalar Quad-Precision Square Root [\& round to Odd] \\
\hline xssubqp[o] & VSX Scalar Quad-Precision Subtract [\& round to Odd] \\
\hline xststdcdp & VSX Scalar Double-Precision Test Data Class \\
\hline xststdcqp & VSX Scalar Quad-Precision Test Data Class \\
\hline xststdcsp & VSX Scalar Single-Precision Test Data Class \\
\hline xsxexpdp & VSX Scalar Double-Precision Extract Exponent \\
\hline xsxexpqp & VSX Scalar Quad-Precision Extract Exponent \\
\hline xsxsigdp & VSX Scalar Double-Precision Extract Significand \\
\hline xsxsigqp & VSX Scalar Quad-Precision Extract Significand \\
\hline xvcvhpsp & VSX Vector Convert Half-Precision to Single-Precision \\
\hline xvcvsphp & VSX Vector round \& Convert Single-Precision to Half-Precision \\
\hline xviexpdp & VSX Vector Double-Precision Insert Exponent \\
\hline xviexpsp & VSX Vector Single-Precision Insert Exponent \\
\hline xvtstdcdp & VSX Vector Double-Precision Test Data Class \\
\hline xvtstdcsp & VSX Vector Single-Precision Test Data Class \\
\hline xvxexpdp & VSX Vector Double-Precision Extract Exponent \\
\hline xvxexpsp & VSX Vector Single-Precision Extract Exponent \\
\hline xvxsigdp & VSX Vector Double-Precision Extract Significand \\
\hline xvxsigsp & VSX Vector Single-Precision Extract Significand \\
\hline xxbrd & VSX Vector Byte-Reverse Dword \\
\hline xxbrh & VSX Vector Byte-Reverse Hword \\
\hline xxbrq & VSX Vector Byte-Reverse Qword \\
\hline xxbrw & VSX Vector Byte-Reverse Word \\
\hline xxextractuw & VSX Vector Extract Unsigned Word \\
\hline xxinsertw & VSX Vector Insert Word \\
\hline xxperm & VSX Vector Permute \\
\hline xxpermr & VSX Vector Permute Right-indexed \\
\hline xxspltib & VSX Vector Splat Immediate Byte \\
\hline
\end{tabular}

\section*{Table 1: Instructions Controlled by the V 2.07 Bit}

Version 2.06 (v2.06)
This bit controls the availability, in problem state, of the following instructions, facilities, and behaviors that were newly available in problem state in the version of the architecture subsequent to Version 2.06.
- icbt
- Iq, stq Ibarx, Iharx, stbcx, sthcx
- Iqarx., stqcx.
- clrbhrb, mfbhrbe
- rfebb, bctar[I]
- The entire Transactional Memory facility
- The instructions in Table 2
- The reserved no-op instructions (see Section 1.9.3 of Book I)
- The reserved SPRs (see Section 1.3.3 of Book I)
- PPR32
- DSCR at SPR number 3
- SIER and MMCR2
- MMCRO \(_{42: 47,51: 55}\) and \(\mathrm{MMCRA}_{0: 63}\).

> Programming Note
> The specified bits of MMCR0 and MMCRA above cannot be changed by mtspr instructions and mfspr instructions return Os for these bits.
- BESCR, EBBHR, and TAR
- The ability of the or 31,31,31 and or \(5,5,5\) instructions to change the value of PPR \({ }_{\text {PRII }}\).
- The ability of mtspr instructions that attempt to set \(\mathrm{PPR}_{\mathrm{PRI}}\) to 001 or 101 to change the value of PPR \(_{\text {PRII }}\).

0 The instructions, facilities, and behaviors listed above are available in problem state.
1 The listed instructions, facilities, and behaviors listed above are unavailable in problem state.
If this bit is set to 1 , then the V 2.07 bit must also be set to 1 .
\begin{tabular}{|c|c|}
\hline Mnemonic & Instruction Name \\
\hline bcdadd. & Decimal Add Modulo \\
\hline bcdsub. & Decimal Subtract Modulo \\
\hline fmrgew & Floating Merge Even Word \\
\hline fmrgow & Floating Merge Odd Word \\
\hline Ixsiwax & Load VSX Scalar as Integer Word Algebraic Indexed \\
\hline Ixsiwzx & Load VSX Scalar as Integer Word and Zero Indexed \\
\hline Ixsspx & Load VSX Scalar Single-Precision Indexed \\
\hline mfvsrd & Move From VSR Doubleword \\
\hline mfvsrwz & Move From VSR Word and Zero \\
\hline mtvsrd & Move To VSR Doubleword \\
\hline mtvsrwa & Move To VSR Word Algebraic \\
\hline mtvsrwz & Move To VSR Word and Zero \\
\hline stxsiwx & Store VSX Scalar as Integer Word Indexed \\
\hline stxsspx & Store VSX Scalar Single-Precision Indexed \\
\hline vaddcuq & Vector Add \& write Carry Unsigned Quadword \\
\hline vaddecuq & Vector Add Extended \& write Carry Unsigned Quadword \\
\hline vaddeuqm & Vector Add Extended Unsigned Quadword Modulo \\
\hline vaddudm & Vector Add Unsigned Doubleword Modulo \\
\hline vadduqm & Vector Add Unsigned Quadword Modulo \\
\hline vbpermq & Vector Bit Permute Quadword \\
\hline vcipher & Vector AES Cipher \\
\hline vcipherlast & Vector AES Cipher Last \\
\hline vclzb & Vector Count Leading Zeros Byte \\
\hline vclzd & Vector Count Leading Zeros Doubleword \\
\hline vclzh & Vector Count Leading Zeros Halfword \\
\hline vclzw & Vector Count Leading Zeros Word \\
\hline vcmpequd[.] & Vector Compare Equal To Unsigned Doubleword \\
\hline vcmpgtsd[.] & Vector Compare Greater Than Signed Doubleword \\
\hline vcmpgtud[.] & Vector Compare Greater Than Unsigned Doubleword \\
\hline veqv & Vector Logical Equivalence \\
\hline vgbbd & Vector Gather Bits by Bytes by Doubleword \\
\hline vmaxsd & Vector Maximum Signed Doubleword \\
\hline vmaxud & Vector Maximum Unsigned Doubleword \\
\hline vminsd & Vector Minimum Signed Doubleword \\
\hline vminud & Vector Minimum Unsigned Doubleword \\
\hline vmrgew & Vector Merge Even Word \\
\hline vmrgow & Vector Merge Odd Word \\
\hline vmulesw & Vector Multiply Even Signed Word \\
\hline vmuleuw & Vector Multiply Even Unsigned Word \\
\hline vmulosw & Vector Multiply Odd Signed Word \\
\hline vmulouw & Vector Multiply Odd Unsigned Word \\
\hline vmuluwm & Vector Multiply Unsigned Word Modulo \\
\hline vnand & Vector Logical NAND \\
\hline
\end{tabular}

Table 2: VSX and Vector Instructions Controlled by the v2.06 Bit
\begin{tabular}{|c|c|}
\hline Mnemonic & Instruction Name \\
\hline vncipher & Vector AES Inverse Cipher \\
\hline vncipherlast & Vector AES Inverse Cipher Last \\
\hline vorc & Vector Logical OR with Complement \\
\hline vpermxor & Vector Permute and Exclusive-OR \\
\hline vpksdss & Vector Pack Signed Doubleword Signed Saturate \\
\hline vpksdus & Vector Pack Signed Doubleword Unsigned Saturate \\
\hline vpkudum & Vector Pack Unsigned Doubleword Unsigned Modulo \\
\hline vpkudus & Vector Pack Unsigned Doubleword Unsigned Saturate \\
\hline vpmsumb & Vector Polynomial Multiply-Sum Byte \\
\hline vpmsumd & Vector Polynomial Multiply-Sum Doubleword \\
\hline vpmsumh & Vector Polynomial Multiply-Sum Halfword \\
\hline vpmsumw & Vector Polynomial Multiply-Sum Word \\
\hline vpopentb & Vector Population Count Byte \\
\hline vpopentd & Vector Population Count Doubleword \\
\hline vpopenth & Vector Population Count Halfword \\
\hline vpopentw & Vector Population Count Word \\
\hline vrld & Vector Rotate Left Doubleword \\
\hline vsbox & Vector AES S-Box \\
\hline vshasigmad & Vector SHA-512 Sigma Doubleword \\
\hline vshasigmaw & Vector SHA-256 Sigma Word \\
\hline vsld & Vector Shift Left Doubleword \\
\hline vsrad & Vector Shift Right Algebraic Doubleword \\
\hline vsrd & Vector Shift Right Doubleword \\
\hline vsubcuq & Vector Subtract \& write Carry Unsigned Quadword \\
\hline vsubecuq & Vector Subtract Extended \& write Carry Unsigned Quadword \\
\hline vsubeuqm & Vector Subtract Extended Unsigned Quadword Modulo \\
\hline vsubudm & Vector Subtract Unsigned Doubleword Modulo \\
\hline vsubuqm & Vector Subtract Unsigned Quadword Modulo \\
\hline vupkhsw & Vector Unpack High Signed Word \\
\hline vupklsw & Vector Unpack Low Signed Word \\
\hline xsaddsp & VSX Scalar Add Single-Precision \\
\hline xscvdpspn & Scalar Convert Double-Precision to Single-Precision format Non-signalling \\
\hline xscvdpspn & Scalar Convert Single-Precision to Double-Precision format Non-signalling \\
\hline xscvsxdsp & VSX Scalar Convert Signed Fixed-Point Doubleword to Single-Precision \\
\hline xscvsxdsp & VSX Scalar round and Convert Signed Fixed-Point Doubleword to Single-Precision format \\
\hline xscvuxdsp & VSX Scalar Convert Unsigned Fixed-Point Doubleword to Single-Precision \\
\hline xscvuxdsp & VSX Scalar round and Convert Unsigned Fixed-Point Doubleword to Single-Precision format \\
\hline xsdivsp & VSX Scalar Divide Single-Precision \\
\hline xsmaddasp & VSX Scalar Multiply-Add Type-A Single-Precision \\
\hline xsmaddmsp & VSX Scalar Multiply-Add Type-M Single-Precision \\
\hline xsmsubasp & VSX Scalar Multiply-Subtract Type-A Single-Precision \\
\hline xsmsubmsp & VSX Scalar Multiply-Subtract Type-M Single-Precision \\
\hline xsmulsp & VSX Scalar Multiply Single-Precision \\
\hline
\end{tabular}

Table 2: VSX and Vector Instructions Controlled by the v2.06 Bit

\section*{Version 3.0}
\begin{tabular}{|l|l|}
\hline Mnemonic & Instruction Name \\
\hline xsnmaddasp & VSX Scalar Negative Multiply-Add Type-A Single-Precision \\
\hline xsnmaddmsp & VSX Scalar Negative Multiply-Add Type-M Single-Precision \\
\hline xsnmsubasp & VSX Scalar Negative Multiply-Subtract Type-A Single-Precision \\
\hline xsnmsubmsp & VSX Scalar Negative Multiply-Subtract Type-M Single-Precision \\
\hline xsresp & VSX Scalar Reciprocal Estimate Single-Precision \\
\hline xsrsp & VSX Scalar Round to Single-Precision \\
\hline xsrsqrtesp & VSX Scalar Reciprocal Square Root Estimate Single-Precision \\
\hline xssqrtsp & VSX Scalar Square Root Single-Precision \\
\hline xssubsp & VSX Scalar Subtract Single-Precision \\
\hline xxleqv & VSX Logical Equivalence \\
\hline xxInand & VSX Logical NAND \\
\hline xxlorc & VSX Logical OR with Complement \\
\hline
\end{tabular}

Table 2: VSX and Vector Instructions Controlled by the v2.06 Bit

62 Version 2.05 (v2.05)
This bit controls the availability, in problem state, of the following instructions, facilities, and behaviors that were newly available in problem state in the version of the architecture subsequent to Version 2.05.
- AMR access using SPR 13
- addg6s
- bperm
- cdtbcd, cbcdtd
- dcffix[.]
- divde[o][.], divdeu[o][.], divwe[o][.], divweu[o][.]
- isel
- Ifiwzx
- fctidu[.], fctiduz[.], fctiwu[.], fctiwuz[.], fcfids[.], fcfidu[.], fcfidus[.], ftdiv, ftsqrt
- Idbrx, stdbrx
- popentw, popcntd
- All facilities in the VSX facility

0 The instructions, facilities, and behaviors listed above are available in problem state.
1 The instructions, facilities, and behaviors listed above are unavailable in problem state.

If this bit is set to 1 , then the v2.06 bit must also be set to 1 .
63 Reserved
The initial state of the PCR is all 0 s.

\section*{Programming Note}

Because the PCR has no effect on privileged instructions except as specified above, privileged instructions that are available on newer implementations but not available on older implementations will behave differently when the thread is in problem state. On older implementations, either an Illegal Instruction type Program interrupt or a Hypervisor Emulation Assistance interrupt will occur because the instruction is undefined; on newer implementations, a Privileged Instruction type Program interrupt will occur because the instruction is implemented. (On older implementations the interrupt will be an Illegal Instruction type Program interrupt if the implementation complies with a version of the architecture that precedes V . 2.05, or complies with V. 2.05 and does not support the Hypervisor Emulation Assistance interrupt, and will be a Hypervisor Emulation Assistance interrupt otherwise.)

In future versions of the architecture, in general the lowest-order reserved bit of the PCR will be used to control the availability of the instructions and related resources that are new in that version of the architecture; the name of the bit will correspond to the previous version of the architecture (i.e., the newest version in which the instructions and related resources were not available).

In these future versions of the architecture, there will be a requirement that if any bit of the low-order defined bits is set to 1 then all higher-order bits of the defined low-order bits must also be set to 1, and the architecture version with which the implementation appears to comply, in problem state, will be the version corresponding to the name of the lowest-order 1 bit in the set of defined low-order PCR bits, or the current architecture version if none of these bits are 1. Also, in general the high-est-order reserved bits will be used to control the availability of sets of instructions and related resources having the requirement that their availability be independent of versions of the architecture.

\subsection*{2.6 Other Hypervisor Resources}

In addition to the resources described above, all hypervisor privileged instructions as well as the following resources are hypervisor resources, accessible to software only when the thread is in hypervisor state except as noted below.

■ All implementation-specific resources except for privileged non-hypervisor implementation-specific SPRs. (See Section 4.4.5 for the list of the imple-mentation-specific SPRs that are allowed to be privileged non-hypervisor SPRs.) Implementa-
tion-specific registers include registers (e.g., "HID" registers) that control hardware functions or affect the results of instruction execution. Examples include resources that disable caches, disable hardware error detection, set breakpoints, control power management, or significantly affect performance.
- ME bit of the MSR

■ SPRs defined as hypervisor-privileged in Section 4.4.5. (Note: Although the Time Base, the PURR, and the SPURR can be altered only by a hypervisor program, the Time Base can be read by all programs and the PURR and SPURR can be read when the thread is in privileged state.)

The contents of a hypervisor resource can be modified by the execution of an instruction (e.g., mtspr) only in hypervisor state \(\left(\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR}=0 \mathrm{~b} 10\right)\). An attempt to modify the contents of a given hypervisor resource, other than \(M S R_{\text {ME }}\), in privileged but non-hypervisor state \(\left(\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR}=0 \mathrm{~b} 00\right)\) causes a Hypervisor Emulation Assistance interrupt when \(\operatorname{LPCR}_{\text {EVIRT }}=1\) and a Privileged Instruction type Program interrupt when \(\mathrm{LPCR}_{\mathrm{EVIRT}}=0\). An attempt to modify \(\mathrm{MSR}_{\text {ME }}\) in privileged but non-hypervisor state is ignored (i.e., the bit is not changed).

\section*{Programming Note}

Because the SPRs listed above are privileged for writing, an attempt to modify the contents of any of these SPRs in problem state ( \(\mathrm{MSR}_{P R}=1\) ) using mtspr causes a Privileged Instruction type Program exception, and similarly for \(\mathrm{MSR}_{\mathrm{ME}}\).

\subsection*{2.7 Sharing Hypervisor Resources}

Shared SPRs are SPRs that are accessible to multiple threads. Changes to shared SPRs made by one thread are immediately readable (using mfspr) by all other threads sharing the SPR.

The LPIDR and DPDES must appear to software to be shared among threads of a sub-processor (see Section 2.8). If the implementation does not support sub-processors, the LPIDR and DPDES must be shared | among all threads of the multi-threaded processor.

Certain additional hypervisor resources may be shared among threads. Programs that modify these resources must be aware of this sharing, and must allow for the fact that changes to these resources may affect more than one thread.

The following additional resources may be shared among threads.
- HRMOR (see Section 2.3)
- LPIDR (see Section 2.4)

■ PCR (see Section 2.5)

■ PVR (see Section 4.3.1)
- RPR (see Section 4.3.8)

I
- PTCR (see Section 5.7.6.1)
- AMOR (see Section 5.7.14.1)

■ HMEER (see Section 6.2.10)
- Time Base (see Section 7.2)
- Virtual Time Base (see Section 7.3)
- Hypervisor Decrementer (see Section 7.5)
- certain implementation-specific registers or imple-mentation-specific fields in architected registers
The set of resources that are shared is implementa-tion-dependent.

Threads that share any of the resources listed above, with the exception of the PTCR, the PVR and the HRMOR, must be in the same partition.
\| For each field of the LPCR, except the AIL, ONL, LD, HDICE, and MER fields, software must ensure that the contents of the field are identical among all threads that are in the same partition and are in a state such that the contents of the field could have side effects. (E.g., software must ensure that the contents of LPCR are identical among all threads that are in the same partition and are not in hypervisor state.) For the HDICE field, software must ensure that the contents of the field are identical among all threads that share the Hypervisor Decrementer and are in a state such that the contents of the field could have side effects. There are no identity requirements for the other fields listed in the first sentence of this paragraph.

\subsection*{2.8 Sub-Processors}

Hardware is allowed to sub-divide a multi-threaded processor into "sub-processors" that appear to privileged programs as multi-threaded processors with fewer threads. Such a multi-threaded processor appears to the hypervisor as a processor with a number of threads equal to the sum of all sub-processor threads, and in which the LPIDR for each sub-processor must appear to be shared among all threads of that sub-processor.

\subsection*{2.9 Thread Identification Register (TIR)}

The TIR is a 64-bit read-only register that contains the thread number, which is a binary number corresponding to the thread.
For implementations that do not support sub-processors, the thread number of a thread is unique among all thread numbers of threads on the multi-threaded processor.

For implementations that support sub-processors, the value of this register depends on whether it is read in hypervisor or privileged, non-hypervisor state as follows.
- When this register is read in privileged, non-hypervisor state, the thread number is unique among all thread numbers of threads on the sub-processor.
- When this register is read in hypervisor state, the thread number is unique among all thread numbers of threads on the multi-threaded processor.

Threads are numbered sequentially, with valid values ranging from 0 to \(t-1\), where \(t\) is the number of threads implemented. A thread for which TIR \(=\mathrm{n}\) is referred to as "thread n."

The layout of the TIR is shown below.


Figure 4. Thread Identification Register
Access to the TIR is privileged.
Since the thread number contained in this register is different if it is read in hypervisor from when it is read in privileged, non-hypervisor state in implementations that support sub-processors, the following conventions are used.
- The value returned in privileged, non-hypervisor state is referred to as the "privileged thread number."
- The value returned in hypervisor state is referred to as the "hypervisor thread number."

\subsection*{2.10 Hypervisor Interrupt Lit-tle-Endian (HILE) Bit}

The Hypervisor Interrupt Little-Endian (HILE) bit is a bit in an implementation-dependent register or similar mechanism. The contents of the HILE bit are copied into \(\mathrm{MSR}_{\text {LE }}\) by interrupts that set \(\mathrm{MSR}_{\mathrm{HV}}\) to 1 (see Section 6.5), to establish the Endian mode for the interrupt handler. The HILE bit is set, by an implementa-tion-dependent method, only during system initialization.
The contents of the HILE bit must be the same for all threads under the control of a given instance of the hypervisor; otherwise all results are undefined.

\section*{Chapter 3. Branch Facility}

\subsection*{3.1 Branch Facility Overview}

This chapter describes the details concerning the registers and the privileged instructions implemented in the Branch Facility that are not covered in Book I.

\subsection*{3.2 Branch Facility Registers}

\subsection*{3.2.1 Machine State Register}

The Machine State Register (MSR) is a 64-bit register. This register defines the state of the thread. On interrupt, the MSR bits are altered in accordance with Figure 64 on page 1061. The MSR can also be modified by the \(\boldsymbol{m t m s r}[d], \boldsymbol{r f s c v}, \boldsymbol{r f i d}\), and \(\boldsymbol{h r f i d}\) instructions. It can be read by the mfmsr instruction.


Figure 5. Machine State Register
Below are shown the bit definitions for the Machine State Register.

Bit Description
\(0 \quad\) Sixty-Four-Bit Mode (SF)
0 The thread is in 32-bit mode.
1 The thread is in 64-bit mode.
1:2 Reserved
\(3 \quad\) Hypervisor State (HV)
0 The thread is not in hypervisor state.
1 If \(\mathrm{MSR}_{\mathrm{PR}}=0\) the thread is in hypervisor state; otherwise the thread is not in hypervisor state.

\section*{Programming Note}

The privilege state of the thread is determined by \(M S R_{H V}\) and MSR \(R_{P R}\), as follows.
\begin{tabular}{ccl} 
HV & PR & \\
0 & 0 & privileged \\
0 & 1 & problem \\
1 & 0 & hypervisor \\
1 & 1 & problem
\end{tabular}

Hypervisor state is also a privileged state \(\left(M_{P R}=0\right)\). All references to "privileged state" in the Books include hypervisor state unless otherwise stated or if it is obvious from the context.
\(\mathrm{MSR}_{\mathrm{HV}}\) can be set to 1 only by the System Call instruction and some interrupts. It can be set to 0 only by rfid and hrfid.
It is possible to run an operating system in an environment that lacks a hypervisor, by always having \(\mathrm{MSR}_{\mathrm{HV}}=1\) and using \(M S R_{H V} \| M_{\text {PR }}=10\) for the operating system (effectively, the OS runs in hypervisor state) and \(M S R_{H V} \| M S R_{P R}=11\) for applications.

\section*{Reserved}

Software must ensure that this bit contains 0; otherwise the results of executing all instructions are boundedly undefined.

\section*{Programming Note}

This bit is initialized to 0 by hardware at system bringup. The handling of this bit by interrupts and by the rfid, hrfid, and \(r f s c v\) instructions is such that, unless software deliberately sets the bit to 1 , the bit will continue to contain 0 .

Reserved
Transaction State (TS)
00 Non-transactional
01 Suspended
10 Transactional
11 Reserved

Changes to MSR[TS] that are caused by Transactional Memory instructions, and by invocation of the transaction's failure handler, take effect immediately (even though these instructions and events are not context synchronizing).
31 Transactional Memory Available (TM)
0 The thread cannot execute any Transactional Memory instructions or access any Transactional Memory registers.
1 The thread can execute Transactional Memory instructions and access Transactional Memory registers unless the Transactional Memory facility has been made unavailable by some other register.

\section*{Vector Available (VEC)}

0 The thread cannot execute any vector instructions, including vector loads, stores, and moves.
1 The thread can execute vector instructions unless they have been made unavailable by some other register.

\section*{Reserved}

\section*{VSX Available (VSX)}

0 The thread cannot execute any VSX instructions, including VSX loads, stores, and moves.
1 The thread can execute VSX instructions unless they have been made unavailable by some other register.

\section*{Programming Note}

An application binary interface defined to support Vector-Scalar operations should also specify a requirement that MSR \(R_{F}\) and \(M_{S R}\) VEC be set to 1 whenever MSR VSX is set to 1.

41:47 Reserved

\section*{External Interrupt Enable (EE)}

0 External, Decrementer, Performance Monitor, and Privileged Doorbell interrupts are disabled.
1 External, Decrementer, Performance Monitor, and Privileged Doorbell interrupts are enabled.

This bit also affects whether Hypervisor Decrementer, Hypervisor Maintenance, and Directed Hypervisor Doorbell interrupts are enabled; see Section 6.5.12 on page 1073, Section 6.5.19 on page 1081, and Section 6.5.20 on page 1081.
Problem State (PR)

0 The thread is in privileged state.
1 The thread is in problem state.

\section*{Programming Note}

Any instruction that sets \(M S R_{P R}\) to 1 also sets \(M S R_{E E}, M S R_{I R}\), and \(M S R_{D R}\) to 1 .

\section*{Floating-Point Available (FP)}

0 The thread cannot execute any float-ing-point instructions, including float-ing-point loads, stores, and moves.
1 The thread can execute floating-point instructions unless they have been made unavailable by some other register.
51 Machine Check Interrupt Enable (ME)
0 Machine Check interrupts are disabled.
1 Machine Check interrupts are enabled.
This bit is a hypervisor resource; see Chapter 2., "Logical Partitioning (LPAR) and Thread Control", on page 927.

\section*{Programming Note}

The only instructions that can alter \(\mathrm{MSR}_{\text {ME }}\) are rfid and hrfid.

00 Trace Disabled: The thread executes instructions normally.
01 Branch Trace: The thread generates a Branch type Trace interrupt after completing the execution of a branch instruction, whether or not the branch is taken.
10 Single Step Trace: The thread generates a Single-Step type Trace interrupt after successfully completing the execution of the next instruction, unless that instruction is an hrfid, rfid, rfscv, or a Power-Saving Mode instruction, all of which are never traced. Successful completion means that the instruction caused no other interrupt and, if the processor is in the Transactional state, is not one of the instructions that is forbidden in Transactional state (e.g., dcbf, see Section 4.3.1 of Book II).

11 Transaction Completion Trace: The thread generates a Transaction Completion type Trace interrupt after the completion of a transaction, whether or not the transaction was successful.
Branch tracing need not be supported. If the function is not implemented, the Ob01 bit encoding is treated as reserved.

Floating-Point Exception Mode 1 (FE1)
See below.
Reserved
Instruction Relocate (IR)
0 Instruction address translation is disabled.
1 Instruction address translation is enabled.

\section*{Programming Note}

See the Programming Note in the definition of MSR \({ }_{P R}\).

Data Relocate (DR)
0 Data address translation is disabled. Effective Address Overflow (EAO) (see Book I) does not occur.
1 Data address translation is enabled. EAO causes a Data Storage interrupt.

\section*{Programming Note}

See the Programming Note in the definition of MSR \({ }_{P R}\).

Reserved

\section*{Performance Monitor Mark (PMM)}

This bit is used by software in conjunction with the Performance Monitor, as described in Chapter 9.

\section*{Programming Note}

Software can use this bit as a pro-cess-specific marker which, in conjunction with MMCR0 \(_{\text {FCM0 FCM1 }}\) (see Section 9.4.4) and MMCR2 (see Section 9.4.6), permits events to be counted on a process-specific basis. (The bit is saved by interrupts and restored by rfid.)

Common uses of the PMM bit include the following.
- All counters count events for a few selected processes. This use requires the following bit settings.
- \(\quad M_{\text {PMM }}=1\) for the selected processes, MSR \(_{\text {PMM }}=0\) for all other processes
- MMCRO \(_{\text {FCMO }}=1\)
- \(\mathrm{MMCRO}_{\mathrm{FCM} 1}=0\)
- \(\quad\) MMCR2 \(=0 \times 0000\)
- All counters count events for all but a few selected processes. This use requires the following bit settings.
- \(\quad M_{\text {PMM }}=1\) for the selected processes, MSR \(_{\text {PMM }}=0\) for all other processes
- \(\mathrm{MMCRO}_{\text {FCMO }}=0\)
- \(\mathrm{MMCRO}_{\mathrm{FCM} 1}=1\)
- \(\quad\) MMCR2 \(=0 \times 0000\)

Notice that for both of these uses a mark value of 1 identifies the "few" processes and a mark value of 0 identifies the remaining "many" processes. Because the PMM bit is set to 0 when an interrupt occurs (see Figure 64 on page 1061), interrupt handlers are treated as one of the "many". If it is desired to treat interrupt handlers as one of the "few", the mark value convention just described would be reversed.

If only a specific counter n is to be frozen, MMCRO \(_{\text {FCM }}\) FCM1 is set to \(0 b 00\), and MMCR2 \(_{\text {FCnM0 }}\) and MMCR2 \({ }_{\text {FCnM1 }}\) instead of MMCRO \({ }_{\text {FCM }}\) and MMCRO \(0_{\text {FCM }}\) are set to the values described above.

\section*{Recoverable Interrupt (RI)}

0 Interrupt is not recoverable.
1 Interrupt is recoverable.
Additional information about the use of this bit is given in Sections 6.4.3, "Interrupt Processing" on page 1057, 6.5.1, "System Reset Interrupt" on page 1062, and 6.5.2, "Machine Check Interrupt" on page 1064

0 The thread is in Big-Endian mode.
1 The thread is in Little-Endian mode.

\section*{Programming Note}

The only instructions that can alter \(\mathrm{MSR}_{\mathrm{LE}}\) are rfid and hrfid.

The Floating-Point Exception Mode bits FE0 and FE1 are interpreted as shown below. For further details see Book I.
\begin{tabular}{ccl} 
FE0 & FE1 & Mode \\
0 & 0 & Ignore Exceptions \\
0 & 1 & Imprecise Nonrecoverable \\
1 & 0 & Imprecise Recoverable \\
1 & 1 & Precise
\end{tabular}

\subsection*{3.2.2 State Transitions Associated with the Transactional Memory Facility}

Updates to \(\mathrm{MSR}_{\text {TS }}\) and \(\mathrm{MSR}_{\text {TM }}\) caused by rfebb, rfid, rfscv, hrfid, or mtmsrd occur as described in Table 3. The value written, and whether or not the instruction causes an interrupt, are dependent on the current values of \(M S R_{T S}\) and \(M S R_{T M}\), and the values being written to these fields. When the setting of \(M S R_{T S}\) causes an illegal state transition, a TM Bad Thing type Program interrupt is generated.

\section*{Programming Note}

The transition rules are the same for mtmsrd as for the rfid-type instructions because if a transition were illegal for mtmsrd but allowed for rfid, or vice versa, software could use the instruction for which the transition is allowed to achieve the effect of the other instruction.

Table 3 shows all the Transaction State transitions that can be requested by \(\boldsymbol{r f e b b}, \boldsymbol{r f i d}, \boldsymbol{r f s c v}, \boldsymbol{h r f i d}\), and \(\boldsymbol{m t m}\) srd. The table covers behavior when TM is enabled by the PCR. For causes of the TM Bad Thing type Program interrupt when TM is disabled by the PCR, see Section 6.5.9. In the table, the contents of \(\mathrm{MSR}_{\text {TS }}\) and \(\mathrm{MSR}_{\mathrm{TM}}\) are abbreviated in the form AB , where A represents \(M S R_{T S}\left(N, T\right.\) or \(S\) ) and \(B\) represents \(M S R_{T M}\) ( 0 or 1). " \(x\) " in the " \(B\) " position means that the entry covers both \(\mathrm{MSR}_{\text {TM }}\) values, with the same value applying in all columns of a given row for a given instance of the transition. (E.g., the first row means that the transition from NO to NO is allowed and results in NO, and that the transition from N0 to N 1 is allowed and results in N1.) "Input \(\mathrm{MSR}_{\mathrm{TS}} \mathrm{MSR}_{\text {TM }}\) " in the second column refers to the \(M S R_{T S}\) and \(M S R_{T M}\) values supplied by CTR for \| rfscv, BESCR for rfebb (just the TS value), SRR1 for rfid, HSRR1 for hrfid, or register RS for mtmsrd.
\begin{tabular}{|c|c|c|c|}
\hline Current
\[
\mathrm{MSR}_{\mathrm{TS}} \mathrm{MSR}_{\mathrm{TM}}
\] & Input
\[
\mathrm{MSR}_{\mathrm{TS}} \mathrm{MSR}_{\mathrm{TM}}
\] & Resulting MSR \(_{\text {TS }}\) MSR \(_{\text {TM }}\) & Comments \\
\hline \multirow[t]{2}{*}{N0} & Nx & Nx & May occur in the context of a Transactional Memory type of Facility Unavailable interrupt handler, enabling/disabling transactions for user-level applications. \\
\hline & All others - Illegal \({ }^{1}\) & NO & \\
\hline т0 & \multicolumn{2}{|c|}{N/A} & Unreachable state \\
\hline \multirow[t]{4}{*}{S0} & N0 \({ }^{2}\) & S0 & Operating system code that is not TM aware may attempt to set TS and TM to zero, thinking they're reserved bits. Change is suppressed. \\
\hline & T1 & T1 & May occur at an rfid returning to an application whose transaction was suspended on interrupt. \\
\hline & Sx & Sx & This case may occur for an rfid returning to an application whose suspended transaction was interrupted. \\
\hline & All others - Illegal \({ }^{1}\) & So & \\
\hline \multirow[t]{2}{*}{N1} & Nx & Nx & After a treclaim, the OS dispatches Nx program. \\
\hline & All others -IIllegal \({ }^{1}\) & N0 & \\
\hline T1 & all & N1 & Disallowed instructions in Transactional state \\
\hline \multirow[t]{3}{*}{S1} & T1 & T1 & \multirow[t]{2}{*}{May occur after trechkpt. when returning to an application.} \\
\hline & Sx & Sx & \\
\hline & All others - Illegal \({ }^{1}\) & So & \\
\hline \multicolumn{4}{|l|}{\begin{tabular}{l}
Notes: \\
1.Generate TM Bad Thing type Program interrupt. "All others" includes all attempts to set MSR \(_{\text {TS }}\) to 0 b11 (reserved value). \\
2. Instruction completes, change to MSR \(_{T M}\) suppressed, except when attempted by rfebb, in which case the result is a TM Bad Thing type Program interrupt.
\end{tabular}} \\
\hline
\end{tabular}

Table 3: Transaction state transitions that can be | requested by rfebb, rfid, rfscv, hrfid, and mtmsrd.

\section*{Version 3.0}

\section*{Programming Note}

For rfscv, [h]rfid, and mtmsrd, the attempted transition from SO to NO is suppressed in order that interrupt handlers that are "unaware" of transactional memory, and load an MSR value that has not been updated to take account of transactional memory, will continue to work correctly. (If the interrupt occurs when a transaction is running or suspended, the interrupt will set MSR[TS II TM] to SO. If the interrupt handler attempts to load an MSR value that has not been updated to take account of transactional memory, that MSR value will have TS II TM = NO. It is desirable that the interrupt handler remain in state S0, so that it can return normally to the interrupted transaction.)
The problem solved by suppressing this transition does not apply to rfebb, so for rfebb an attempt to transition from SO to NO is not suppressed, and instead causes a TM Bad Thing type Program interrupt.

\subsection*{3.2.3 Processor Stop Status and Control Register (PSSCR)}

The layout of the PSSCR is shown below.


Figure 6. Processor stop Status and Control Register

The contents of the PSSCR control the operation of the stop instruction and provide status indicating the level of power saving that was entered while in power-saving mode.
All fields of this register can be read and written by the hypervisor using either hypervisor SPR 855 or privileged SPR 823. A subset of the fields of this register can be read and written in privileged state using privileged SPR 823, as specified below. Fields that can only be read or written by the hypervisor are indicated below; all other fields can be read or written in either privileged or hypervisor states. When a field that is accessible only to the hypervisor is accessed in privileged state, writes have no effect and reads return Os regardless of the value of the field.
The bits and their meanings are as follows.

\section*{0:3 Power-Saving Level Status (PLS) \\ Hardware sets this field to the highest power-saving level that the thread entered between the time when the stop instruction is executed and when the thread exits power-saving mode. See the description of the SD field for the value returned in this field when the PSSCR is read.}

\section*{- Programming Note}

Since the power-saving level entered during power-saving mode may vary with time, the PLS field may not indicate the power-saving level that existed at exit from power-saving mode.

\section*{4:40 Reserved}
\(41 \quad\) Status Disable (SD)
This field is accessible only to the hypervisor.
0 The current value of the PLS field is returned in the PLS field when reading the PSSCR (using mfspr).
1 O's are returned in the PLS field when reading the PSSCR (using mfspr).

\section*{Programming Note}

Before dispatching an OS, the hypervisor may initialize this field to 1 in order to prevent the OS from reading the Power-Saving Level Status (PLS) field. This may be necessary in secure environments since an OS may be capable of detecting the presence of another OS on the same processor by observing the state of the PLS field after exiting power-saving mode.

\section*{Enable State Loss (ESL)}

This field is accessible only to the hypervisor.
0 State loss while in power-saving mode is controlled by the RL, MTL, and PSLL fields.
1 Non-hypervisor state loss is allowed while in power-saving mode in addition to state loss controlled by the RL, MTL, and PSLL fields.

If this field is set to 1 when the stop instruction is executed in privileged non-hypervisor state, a Hypervisor Facility Unavailable interrupt occurs. See Section 6.5.26.

\section*{Programming Note}

When state loss occurs, thread resources such as SPRs, GPRs, address translation resources, etc. may be powered off or allocated to other threads during power-saving mode. The amount of state loss for various combinations of ESL, RL, and MTL values is implementation dependent, subject to the restrictions specified in Section 3.3.2.

\section*{Exit Criterion (EC)}

This field is accessible only to the hypervisor.
0 Hardware will exit power-saving mode when the exception corresponding to any system-caused interrupt occurs.
Power-saving mode is exited either at the
instruction following the stop (if \(\mathrm{MSR}_{\mathrm{EE}}=0\) ) or in the corresponding interrupt handler (if \(\mathrm{MSR}_{E E}=1\) ).
1 Provided LPCR \({ }_{\text {PECE }}\) is not lost, hardware will exit power-saving mode only when a System Reset exception or one of the events specified in LPCR PECE occurs. If the event is a Machine Check exception, then a Machine Check interrupt occurs; otherwise a System Reset interrupt occurs, and the contents of SRR1 indicate the event that caused exit from power-saving mode.
When the stop instruction is executed in hypervisor state, the hypervisor must set the ESL field to the same value as this field. Also, if the RL or MTL fields are set to values that allow state loss, then fields ESL and EC must both be set to 1. Other combinations of the values of the ESL, EC, RL, and MTL fields are reserved for future use.

\section*{Architecture Note}

Other combinations of the values of the ESL, EC, RL, and MTL fields may be allowed in a future version of the architecture in order to provide additional functionality.

If this field is set to 1 when the stop instruction is executed in privileged state, a Hypervisor Facility Unavailable interrupt occurs. See Section 6.5.26.

\section*{Programming Note}

In order to enable an OS to enter power-saving mode without hypervisor involvement, both the EC and ESL bits must be set to Os. When this is done, OS execution of the stop instruction will not cause hypervisor involvement provided that bits RL and and MTL are less than or equal to PSLL. See Section 6.5.26 for details.

44:47 Power-Saving Level Limit (PSLL)
This field is accessible only to the hypervisor.
This field limits the power-saving level that may be entered or transitioned into when the stop instruction is executed in privileged, non-hypervisor state; when the stop instruction is executed in hypervisor state, this field is ignored.
Reserved
Transition Rate (TR)

This field is used to specify the relative rate at which the power-saving level increases during power-saving mode. The rate of power-saving level increase corresponding to each value is implementation-dependent, and monotonically increasing with the value specified.
56:59 Maximum Transition Level (MTL)
If the value of this field is greater than the value of the Power-Saving Level Limit (PSLL) field when stop is executed in privileged, non-hypervisor state, a Hypervisor Facility Unavailable interrupt occurs. See Section 6.5.26 of Book III.

Otherwise, if the value of this field is greater than the value of the RL field, the power-saving level is allowed to increase from the value in the RL field up to the value of this field during power-saving mode.

If this field is less than or equal to the value of the PSLL field when stop is executed in privileged non-hypervisor state, this field is used to specify the maximum power-saving level that can be reached during power-saving mode provided that the value of this field is greater than the value of the RL field. If this field is less than the Requested Level (RL) field when stop is executed hardware is not allowed to increase the power-saving level during power-saving mode beyond the value indicated in the RL field.

\section*{Requested Level (RL)}

This field is used to specify the power-saving level that is to be entered when the stop instruction is executed.

If the value of this field is greater than the value of the Power-Saving Level Limit (PSLL) field when stop is executed in privileged, non-hypervisor state, a Hypervisor Facility Unavailable interrupt occurs.

\section*{Programming Note}

The Hypervisor Facility Unavailable interrupt occurs when a privileged non-hypervisor program executes stop when PSSCR \(_{\text {RL }}>\) PSSCR \(_{\text {PSLL }}\) so that the Hypervisor may decide whether or not to allow the requested loss of state to occur.

If the hypervisor decides that some loss of state is acceptable, it may choose to re-execute stop after either setting PSS\(\mathrm{CR}_{\mathrm{MTL}}\) to a value that causes state loss, or setting both PSSCR \(_{\text {RL }}\) and PSSCR MTL to values that cause state loss. When the thread exits power-saving mode, the hypervisor can quickly determine whether any resources were actually lost and need to be restored.

\subsection*{3.3 Branch Facility Instructions}

\subsection*{3.3.1 System Linkage Instructions}

These instructions provide the means by which a program can call upon the system to perform a service, and by which the system can return from performing a service or from processing an interrupt.

The System Call instruction is described in Book I, but only at the level required by an application programmer. A complete description of this instruction appears below.

\section*{System Call}

SC-form

```

SRRO }\mp@subsup{\leftarrow}{iea}{CIA + 4
SRR1 33:36 42:47 }\leftarrow
SRR10:32 37:41 48:63 \leftarrow MSRR0:32 37:41 48:63
MSR }\leftarrow\mathrm{ new_value (see below)
NIA \leftarrow 0x0000_0000_0000_0C00

```

The effective address of the instruction following the System Call instruction is placed into SRR0. Bits 0:32, 37:41, and 48:63 of the MSR are placed into the corresponding bits of SRR1, and bits 33:36 and 42:47 of SRR1 are set to zero.

Then a System Call interrupt is generated. The interrupt causes the MSR to be set as described in Section 6.5, "Interrupt Definitions" on page 1060. The setting of the MSR is affected by the contents of the LEV field. LEV values greater than 1 are reserved. Bits \(0: 5\) of the LEV field (instruction bits 20:25) are treated as a reserved field.

The interrupt causes the next instruction to be fetched from effective address 0x0000_0000_0000_0C00.
This instruction is context synchronizing.

\section*{Special Registers Altered:}

SRRO SRR1 MSR

\section*{Programming Note}

If \(L E V=1\) the hypervisor is invoked. This is the only way that executing an instruction can cause hypervisor state to be entered.

Because this instruction is not privileged, it is possible for application software to invoke the hypervisor. However, such invocation should be considered a programming error.

\section*{Programming Note}
sc serves as both a basic and an extended mnemonic. The Assembler will recognize an sc mnemonic with one operand as the basic form, and an sc mnemonic with no operand as the extended form. In the extended form the LEV operand is omitted and assumed to be 0 .

System Call Vectored
SC-form
scv LEV

```

LR}\leftarrowCIA + 4
CTR 33:36 42:47 }\leftarrow undefined
CTR 0:32 37:41 48:63 \leftarrow MSR0:32 37:41 48:63
MSR \leftarrow new_value (see below)
NIA}\leftarrow\mathrm{ (see below)

```

The effective address of the instruction following the System Call Vectored instruction is placed into the Link Register. Bits 0:32, 37:41, and 48:63 of the MSR are placed into the corresponding bits of Count Register, and bits 33:36 and 42:47 of Count Register are set to undefined values.

Then a System Call Vectored interrupt is generated. The interrupt causes the MSR to be altered as described in Section 6.5.

The interrupt causes the next instruction to be fetched as specified in LPCR \({ }_{\text {AIL }}\) (see to Section 2.2).

The SRRs are not affected.
This instruction is context synchronizing.

\section*{Special Registers Altered:}

LR CTR MSR

Return From System Call Vectored XL-form
rfscv
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline \multicolumn{1}{|c|}{19} & \multicolumn{2}{|c|}{\(/ / /\)} & \multicolumn{2}{|c|}{\(/ / /\)} & \multicolumn{2}{|c|}{\(/ / /\)} \\
0 & & 6 & & 11 & 16 & 21 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { if }\left(\mathrm{MSR}_{29: 31} \text { ᄀ= } 0 \mathrm{~b} 010 \mid \mathrm{CTR}_{29: 31} \text { ㄱ= } 0 \mathrm{~b} 000\right) \text { then } \\
& \mathrm{MSR}_{29: 31} \leftarrow \mathrm{CTR}_{29: 31} \\
& \mathrm{MSR}_{48} \leftarrow \mathrm{CTR}_{48} \mid \mathrm{CTR}_{49} \\
& \begin{array}{l|l}
\mathrm{MSR}_{58} \leftarrow \mathrm{CTR}_{58} & \mathrm{CTR}_{49} \\
\mathrm{MSR}_{59} \leftarrow \mathrm{CTR}_{59} & \mathrm{CTR}_{49}
\end{array} \\
& \mathrm{MSR}_{0: 2} \text { 4:28 } 32 \text { 37:41 49:50 52:57 60:63 } \leftarrow^{\leftarrow} \mathrm{CTR}_{0: 2} 4: 2832 \text { 37:41 49:50 52:57 } \\
& \text { 60:63 } \\
& \text { NIA } \leftarrow_{\text {iea }}{L R_{0: 61 ~}}| | 0000
\end{aligned}
\]

If bits 29 through 31 of the MSR are not equal to 0b010 or bits 29 through 31 of CTR are not equal to 0b000, then the value of bits 29 through 31 of CTR is placed into bits 29 through 31 of the MSR. The result of ORing bits 48 and 49 of CTR is placed into \(\mathrm{MSR}_{48}\). The result of ORing bits 58 and 49 of CTR is placed into \(M S R_{58}\). The result of ORing bits 59 and 49 of CTR is placed into \(\mathrm{MSR}_{59}\). Bits 0:2, 4:28, 32, 37:41, 49:50, 52:57, and 60:63 of CTR are placed into the corresponding bits of the MSR

If the instruction attempts to cause an illegal transaction state transition (see Table 3, "Transaction state transitions that can be requested by rfebb, rfid, rfscv, hrfid, and mtmsrd.," on page 947), or when TM is disabled by the PCR, a transition to Problem state with an active transaction, a TM Bad Thing type Program interrupt is generated (unless a higher-priority exception is pending). If this interrupt is generated, the value placed into SRRO by the interrupt processing mechanism (see Section 6.4.3) is the address of the rfid instruction. Otherwise, if the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address \(\mathrm{LR}_{0: 61}\) II \(0 b 00\) (when \(\mathrm{SF}=1\) in the new MSR value) or \({ }^{32} 0\) II \(\mathrm{LR}_{32: 61}\) II \(0 b 00\) (when \(\mathrm{SF}=0\) in the new MSR value). If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRRO or HSRRO by the interrupt processing mechanism (see Section 6.4.3) is the address of the instruction that would have been executed next had the interrupt not occurred.

This instruction is privileged and context synchronizing.

\section*{Special Registers Altered: MSR}

\section*{Programming Note}

If this instruction sets \(M_{\text {MRR }}\) to 1 , it also sets \(M_{\text {EE }}, \mathrm{MSR}_{\mathrm{IR}}\), and \(\mathrm{MSR}_{\mathrm{DR}}\) to 1 .

\section*{Return From Interrupt Doubleword XL-form}
rfid
\begin{tabular}{|l|l|l|l|l|ll|l|}
\hline \multicolumn{1}{|c|}{19} & \multicolumn{2}{|c|}{\(/ / /\)} & \multicolumn{1}{|c|}{\(/ / /\)} & \multicolumn{2}{|c|}{\(/ / /\)} & & 18 \\
0 & & 6 & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

$\mathrm{MSR}_{51} \leftarrow\left(\mathrm{MSR}_{3} \& \mathrm{SRR}_{51}\right) \mid\left(\left(\neg \mathrm{MSR}_{3}\right) \& \mathrm{MSR}_{51}\right)$
$\mathrm{MSR}_{3} \leftarrow \mathrm{MSR}_{3} \& \mathrm{SRR}_{3}$
if $\left(\mathrm{MSR}_{29: 31} \neg=0 \mathrm{~b} 010 \mid \mathrm{SRR1}_{29: 31}\right.$ ㄱ= 0b000) then
$\mathrm{MSR}_{29: 31} \leftarrow \mathrm{SRR}_{29: 31}$
$\mathrm{MSR}_{48} \leftarrow \mathrm{SRR}_{48} \mid \mathrm{SRR1}_{49}$
$\mathrm{MSR}_{58} \leftarrow \mathrm{SRR}_{58} \mid \mathrm{SRR}_{49}$
MSR $_{59} \leftarrow \operatorname{SRR}_{59} \mid \operatorname{SRR}_{49}$
MSR $_{0}: 2$ 4:28 32 37:41 49:50 52:57 60:63 $\leftarrow$ SRR10:2 4:28 32 37:41 49:50 52:57
60:63
NIA $\leftarrow_{\text {iea }}$ SRRO $_{0: 61}| | 0$ 0.b00

```

If \(\mathrm{MSR}_{3}=1\) then bits 3 and 51 of SRR1 are placed into the corresponding bits of the MSR. If bits 29 through 31 of the MSR are not equal to 0 b 010 or bits 29 through 31 of SRR1 are not equal to 0b000, then the value of bits 29 through 31 of SRR1 is placed into bits 29 through 31 of the MSR. The result of ORing bits 48 and 49 of SRR1 is placed into \(\mathrm{MSR}_{48}\). The result of ORing bits 58 and 49 of SRR1 is placed into \(\mathrm{MSR}_{58}\). The result of ORing bits 59 and 49 of SRR1 is placed into \(\mathrm{MSR}_{59}\). Bits 0:2, 4:28, 32, 37:41, 49:50, 52:57, and 60:63 of SRR1 are placed into the corresponding bits of the MSR.
If the instruction attempts to cause an illegal transaction state transition (see Table 3, "Transaction state transitions that can be requested by rebb, rfid, rfscv, hrfid, and mtmsrd.," on page 947), or when TM is disabled by the PCR, a transition to Problem state with an active transaction, a TM Bad Thing type Program interrupt is generated (unless a higher-priority exception is pending). If this interrupt is generated, the value placed into SRRO by the interrupt processing mechanism (see Section 6.4.3) is the address of the rfid instruction. Otherwise, if the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address \(\mathrm{SRRO}_{0: 61}\) II 0 b 00 (when \(\mathrm{SF}=1\) in the new MSR value) or \({ }^{32} 0\) II \(\mathrm{SRRO}_{32: 61}\) II 0 bOO (when \(\mathrm{SF}=0\) in the new MSR value). If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRRO or HSRRO by the interrupt processing mechanism (see Section 6.4.3) is the address of the instruction that would have been executed next had the interrupt not occurred.

This instruction is privileged and context synchronizing.

\section*{Special Registers Altered: MSR}

\section*{Hypervisor Return From Interrupt Doubleword XL-form}
hrfid
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline \multicolumn{1}{|c|}{19} & \multicolumn{1}{c|}{\(/ / /\)} & \multicolumn{1}{|c|}{\(/ / /\)} & \multicolumn{1}{c|}{274} & \(/\) \\
0 & & 6 & & 11 & 16 & 21 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { if }\left(\mathrm{MSR}_{29: 31} \neg=0 \mathrm{bb010} \mid \operatorname{HSRR}_{29: 31} \neg=0 \mathrm{~b} 000\right) \text { then } \\
& \text { MSR }_{29: 31} \leftarrow \text { HSRR1 }_{29: 31} \\
& \text { MSR }_{48} \leftarrow \mathrm{HSRR1}_{48} \mid \mathrm{HSRR1}_{49} \\
& \text { MSR }_{58} \leftarrow \mathrm{HSRR}_{58} \mid \mathrm{HSRR}_{49} \\
& \text { MSR }_{59} \leftarrow \mathrm{HSRR}_{59} \mid \mathrm{HSRR}_{49} \\
& \text { MSR }_{0: 28} 3237: 4149: 57 \text { 60:63 } \leftarrow \text { HSRR1 }_{0: 28} 32 \text { 37:41 49:57 60:63 } \\
& \text { NIA } \leftarrow_{\text {iea }} \operatorname{HSRRO}_{0: 61} \| 0 \mathrm{Ob0} 0
\end{aligned}
\]

If bits 29 through 31 of the MSR are not equal to 0b010 or bits 29 through 31 of HSRR1 are not equal to 0b000, then the value of bits 29 through 31 of HSRR1 is placed into bits 29 through 31 of the MSR. The result of ORing bits 48 and 49 of HSRR1 is placed into \(\mathrm{MSR}_{48}\). The result of ORing bits 58 and 49 of HSRR1 is placed into \(\mathrm{MSR}_{58}\). The result of ORing bits 59 and 49 of HSRR1 is placed into \(\mathrm{MSR}_{59}\). Bits \(0: 28,32,37: 41,49: 57\), and 60:63 of HSRR1 are placed into the corresponding bits of the MSR.

If the instruction attempts to cause an illegal transaction state transition (see Table 3, "Transaction state transitions that can be requested by rebb, rfid, rfscv, hrfid, and mtmsrd.," on page 947), or when TM is disabled by the PCR, a transition to Problem state with an active transaction, a TM Bad Thing type Program interrupt is generated (unless a higher-priority exception is pending). If this interrupt is generated, the value placed into SRRO by the interrupt processing mechanism (see Section 6.4.3) is the address of the hrfid instruction. Otherwise, if the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address \(\operatorname{HSRRO}_{0: 61}\) II 0 b 00 (when \(\mathrm{SF}=1\) in the new MSR value) or \({ }^{32} 0\) I| \(\mathrm{HSRRO}_{32: 61}\) II \(0 b 00\) (when \(\mathrm{SF}=0\) in the new MSR value). If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRRO or HSRRO by the interrupt processing mechanism (see Section 6.4.3) is the address of the instruction that would have been executed next had the interrupt not occurred.

This instruction is hypervisor privileged and context synchronizing.
Special Registers Altered: MSR

\section*{Programming Note}

If this instruction sets \(M S R_{P R}\) to 1 , it also sets \(M^{\prime} R_{E E}, M S R_{I R}\), and \(M S R_{D R}\) to 1 .

\subsection*{3.3.2 Power-Saving Mode}

Power-Saving Mode is a mode in which the thread does not execute instructions and may consume less power than it would if it were not in power-saving mode.

There are 16 levels of power savings, designated as levels 0-15. For each power-saving level, the power consumed may be less than or equal to the power consumed in the next-lower level, and the time required for the thread to exit power-saving mode and resume execution may be greater than or equal that of the next-lower level.

When the thread is in power-saving mode, some resource state may be lost. The state that may be lost while in each power-saving level is implementation dependent, with the following restrictions.
- For PSSCR \(_{E S L}=0\) and power-saving level 0000, no thread state is lost.
- There must be a power-saving level in which the Decrementer and all hypervisor resources are maintained as if the thread was not in power-saving mode, and in which sufficient information is maintained to allow the hypervisor to resume execution.
- The amount of state loss in a given level is less than or equal to the amount of state loss in the next higher level.
- The state of all read-only resources and the HRMOR is always maintained.

\section*{Programming Note}

For the power-saving level corresponding to the second item above, if the state of the Decrementer were not maintained and updated as if the thread was not in power-saving mode, Decrementer exceptions would not reliably cause exit from this power-saving level even if Decrementer exceptions were enabled to cause exit.

The thread can be put in power-saving mode by executing the stop instruction. As specified below, this instruction stops execution immediately after the stop instruction is executed, and the thread is put into power-saving mode. The power-saving level that is entered depends on the contents of the PSSCR (see Section 3.2.3).

\subsection*{3.3.2.1 Power-Saving Mode Instruction}

The stop instruction is used to stop instruction fetching and execution and put the thread into power-saving mode. The thread remains in power-saving mode until
a system reset exception or an event that is enabled to cause exit from power-saving mode occurs. (See the definition of PSSCR \(_{E C}\) inSection 3.2.3.)


The thread is placed into power-saving mode and execution is stopped.

The power-saving level that is entered is determined by the contents of the PSSCR (see Section 3.2.3). The thread state that is maintained depends on the power-saving level that is entered. The thread state that is maintained at each power-saving level is implemen-tation-dependent, subject to the restrictions specified in Section 3.3.2.

The thread remains in power-saving mode until either a System Reset exception or certain other events occur. The events that may cause exit from power-saving mode are specified by PSSCR EC and LPCR PECE . If the event that causes the exit is a System Reset, Machine Check, or Hypervisor Maintenance exception, resource state that would be lost if the exception occurred when the thread was not in power-saving mode may be lost.

An attempt to execute this instruction in Suspended state will result in a TM Bad Thing type of Program interrupt.

This instruction is privileged and context synchronizing.

\section*{Special Registers Altered: \\ None}

\subsection*{3.3.2.2 Entering and Exiting Power-Saving Mode}

Before software executes the stop instruction, the PSSCR is initialized. If the stop instruction is to be used by the OS, the hypervisor initializes the fields that are accessible only to the hypervisor before dispatching the OS. These fields include the SD, ESL, EC, and PSLL fields. See the Programming Notes for these fieldsSection 3.2.3 for additional information.
If the stop instruction is to be executed by the hypervisor when PSSCR \(_{E C}=1\), the LPCR \(_{\text {PECE }}\) must be set to the desired value (see Section 2.2). Depending on the implementation and the power-saving level to be entered, it may also be necessary to save the state of certain resources and perform synchronization procedures to ensure that all stores have been performed with respect to other threads or mechanisms that use
the storage areas before executing the stop. See the the User's Manual for the implementation for details.
Software must also specify the requested and maximum power-saving level limit fields (i.e RL and MTL fields), and the Transition Rate (TR) field in the PSSCR in order to bound the range of power-saving modes that can be entered. If the value of the RL field is greater than or equal to the value of the MTL field, the power-saving level will not increase from the initial level during power-saving mode.

\section*{Programming Note}

If \(\mathrm{MSR}_{\mathrm{EE}}=1\) when the stop instruction is executed, then the interrupt corresponding to the exception that was expected to cause exit from power-saving mode may occur immediately prior to execution of the stop instruction. If this occurs, the result may be a software hang condition since the exception that was expected to cause exit from power-saving mode has already occurred.

The above software hang condition can be prevented by setting \(\mathrm{MSR}_{\mathrm{EE}}=0\) prior to executing stop.

After the thread has entered power-saving mode with \(\mathrm{PSSCR}_{E C}=0\), any exception may cause exit from power-saving mode. When an exception occurs, power-saving mode is exited either at the instruction following the stop (if \(\mathrm{MSR}_{\mathrm{EE}}=0\) ) or in the corresponding interrupt handler (if \(\mathrm{MSR}_{\mathrm{EE}}=1\) ).

\section*{Programming Note}

If stop was executed when PSSCR \(_{E C}=0\) and \(\mathrm{MSR}_{\mathrm{EE}}=0\) (in order to avoid the hang condition described in the above Programming Note), \(\mathrm{MSR}_{\text {EE }}\) should be set to 1 after power-saving mode is exited in order to take the interrupt corresponding to the exception that caused exit from power-saving mode.

After the thread has entered power-saving mode with PSSCR \(_{E C}=1\), only the System Reset or Machine Check exceptions and the exceptions enabled in LPCR \({ }_{\text {PECE }}\) will cause exit. If the event that causes exit is a Machine Check exception, then a Machine Check interrupt occurs; otherwise a System Reset interrupt occurs, and the contents of SRR1 indicate the exception that caused exit from power-saving mode.
If the hypervisor has set PSSCR \(_{S D}=0\) prior to when the stop instruction is executed, the instruction following the stop may typically be a mfspr in order to read the

Version 3.0
contents of PSSCR PLS to determine the maximum power-saving level that was entered during power-saving mode.

\subsection*{3.4 Event-Based Branch Facility and Instruction}

The Event-Based Branch facility is described in Chapter 7 of Book II, but only at the level required by the application program.

Event-based branches and event-based exceptions can only occur in problem state and when event-based branches and exceptions have been enabled in the FSCR and HFSCR, and BESCR \({ }_{\text {GE }}=1\). Additionally, the following additional bits must be set to one in order to enable EBB exceptions specific to a given function to occur.
- \(\mathrm{MMCRO}_{\text {EBE }}\) and BESCR PME must be set to 1 to enable Performance Monitor event-based exceptions.
- FSCR \(_{\text {LM }}\) and BESCR LME must be set to 1 to enable Load Monitored event-based exceptions.
- \(B_{E S C R}\) EE must be set to 1 to enable External event-based exceptions.

If an event-based exception exists when \(M S R_{P R}=0\), the corresponding event-based branch does not occur until \(\mathrm{MSR}_{\mathrm{PR}}=1\), \(\mathrm{FSCR}_{\mathrm{EBB}}=1\), HFSCR EBB \(=1\), and BES\(\mathrm{CR}_{\mathrm{GE}}=1\).

If the rfebb instruction attempts to cause a transition to Transactional or Suspended state when \(\mathrm{PCR}_{T M}=1\) or an illegal transaction state transition (see Section 3.2.2), a TM Bad Thing type Program interrupt is generated (unless a higher-priority exception is pending). If this interrupt is generated, the value placed into SRRO by the interrupt processing mechanism is the address of the rfebb instruction.)

Version 3.0

\section*{Chapter 4. Fixed-Point Facility}

\subsection*{4.1 Fixed-Point Facility Overview}

This chapter describes the details concerning the registers and the privileged instructions implemented in the Fixed-Point Facility that are not covered in Book I.

\subsection*{4.2 Special Purpose Registers}

Special Purpose Registers (SPRs) are read and written using the mfspr (page 975) and mtspr (page 974) instructions. Most SPRs are defined in other chapters of this book; see the index to locate those definitions.

\subsection*{4.3 Fixed-Point Facility Registers}

\subsection*{4.3.1 Processor Version Register}

The Processor Version Register (PVR) is a 32-bit read-only register that contains a value identifying the version and revision level of the implementation. The contents of the PVR can be copied to a GPR by the mfspr instruction. Read access to the PVR is privileged; write access is not provided.


Figure 7. Processor Version Register
The PVR distinguishes between implementations that differ in attributes that may affect software. It contains two fields.
Version A 16-bit number that identifies the version of the implementation. Different version numbers indicate major differences between implementations.

Revision A 16-bit number that distinguishes between implementations of the version. Different revision numbers indicate minor differences between implementations having the same
version number, such as clock rate and Engineering Change level.

Version numbers are assigned by the Power ISA process. Revision numbers are assigned by an implemen-tation-defined process.

\subsection*{4.3.2 Chip Information Register}

The Chip Information Register (CIR) is a 32-bit read-only register that contains a value identifying the manufacturer and other characteristics of the chip on which the processor is implemented. The contents of the CIR can be copied to a GPR by the mfspr instruction. Read access to the CIR is privileged; write access is not provided.


\section*{Bit Description}

32:35 Manufacturer ID (ID) A four-bit field that identifies the manufacturer of the chip.

36:63 Implementation-dependent.
Figure 8. Chip Information Register

\subsection*{4.3.3 Processor Identification Register}
\(\qquad\)
The Processor Identification Register (PIR) is a 32-bit
I register that contains a 20-bit PROCID field that can be used to distinguish the thread from other threads in the system. The contents of the PIR can be copied to a GPR by the mfspr instruction. Read access to the PIR is privileged; write access is not provided.

\begin{tabular}{lll} 
Bits & Name & Description \\
\(32: 43\) & & Reserved \\
\(44: 63\) & PROCID & Thread ID
\end{tabular}

Figure 9. Processor Identification Register
The means by which the PIR is initialized are imple-mentation-dependent.

The PIR is a hypervisor resource; see Chapter 2.

\subsection*{4.3.4 Process Identification Register}

The layout of the Process Identification Register (PIDR) is shown in Figure 10 below.


\section*{Bit(s) Name Description \\ 32:63 PID Process Identifier}

Figure 10. Process Identification Register
The contents of the PIDR identify the process to which the thread is assigned. The value is used to perform translation and manage the caching of translations. The number of PIDR bits supported is implementa-tion-dependent.

Access to the PIDR is privileged.

\section*{Programming Note}

Radix tree translation assigns special meaning to PID=0, specifically indicating the operating system's kernel process. When GR=1, PIDR should not be set to zero except when \(M_{\text {PR }}=0\).

\subsection*{4.3.5 Control Register}

The Control Register (CTRL) is a 32-bit register as shown below.
\begin{tabular}{|l|r|r|r|}
\hline /// & TS & /// & RUN \\
\hline 32 & 48 & 56 & 63 \\
\hline
\end{tabular}

Figure 11. Control Register
The field definitions for the CTRL are shown below.

\section*{Bit(s) Description}

32:47 Reserved
48:55 Thread State (TS)

Problem State Access
Reserved
Privileged accesses
Bits \(0: 7\) of this field are read-only bits that indicate the state of CTRL \({ }_{\text {RUN }}\) for threads with privileged thread numbers 0 through 7, respectively; bits corresponding to privileged thread numbers higher than the maximum privileged thread number supported are set to Os.

\section*{Hypervisor accesses}

Bits 0:7 of this field are read-only bits that indicate the state of CTRL RUN for threads with hypervisor thread numbers 0 through 7, respectively; bits corresponding to hypervisor thread numbers higher than the maximum hypervisor thread number supported are set to 0 s .

56:62 Reserved
63 RUN
This bit controls an external I/O pin. This signal may be used for the following:
- driving the RUN Light on a system operator panel
■ Direct External exception routing
- Performance Monitor Counter incrementing (see Chapter 9)
The RUN bit can be used by the operating system to indicate when the thread is doing useful work.

Write access to the CTRL is privileged. Reads can be performed in privileged or problem state.

\subsection*{4.3.6 Program Priority Register}

Privileged programs may set a wider range of program priorities in the PRI field of PPR and PPR32 than may be set by problem state programs (see Chapter 3 of Book II). Problem state programs may only set values in the range of \(0 b 001\) to \(0 b 100\) unless the Problem State Priority Boost register (see Section 4.3.7) allows the value 0b101. Privileged programs may set values in the range of 0b001 to 0b110. Hypervisor software may also set 0b111. For all priorities except 0b101, if a program attempts to set a value that is not allowed for its privilege level, the PRI field remains unchanged. If a problem state program attempts to set its priority value to Ob101 when this priority value is not allowed for problem state programs, the priority is set to \(0 b 100\). The values and their corresponding meanings are as follows.

\section*{Bit(s) Description}

11:13 Program Priority (PRI)
001 very low

\section*{010 low \\ 011 medium low \\ 100 medium \\ 101 medium high \\ 110 high \\ 111 very high \\ 4.3.7 Problem State Priority Boost Register}

The Problem State Priority Boost (PSPB) register is a 32-bit register that controls whether problem state programs have access to program priority medium high. (See Section 3.1 of Book II.)


Figure 12. Problem State Priority Boost Register
A problem state program is able to set the program priority to medium high only when the PSPB of the thread contains a non-zero value.

The maximum value to which the PSPB can be set must be a power of 2 minus 1 . Bits that are not required to represent this maximum value must return 0s when read regardless of what was written to them.

When the PSPB is set to a value less than its maximum value but greater than 0 , its contents decrease monotonically at the same rate as the SPURR until its contents minus the amount it is to be decreased are 0 or less when a problem state program is executing on the thread at a priority of medium high. When the contents of the PSPB minus the amount it is to be decreased are 0 or less, its contents are replaced by 0.

When the PSPB is set to its maximum value or 0 , its contents do not change until it is set to a different value.

Whenever the priority of a thread is medium high and either of the following conditions exist, hardware changes the priority to medium:
- the PSPB counts down to 0 , or
- \(\quad \mathrm{PSPB}=0\) and the privilege state of the thread is changed to problem state \(\left(\mathrm{MSR}_{P R}=1\right)\).

\subsection*{4.3.8 Relative Priority Register}

The Relative Priority Register (RPR) is a 64-bit register that allows the hypervisor to control the relative priorities corresponding to each valid value of PPR \(_{\text {PRII }}\).
\begin{tabular}{|l|l|l|l|l|l|l|l|}
\hline\(/\) & \(\mathrm{RP}_{1}\) & \(\mathrm{RP}_{2}\) & \(\mathrm{RP}_{3}\) & \(\mathrm{RP}_{4}\) & \(\mathrm{RP}_{5}\) & \(\mathrm{RP}_{6}\) & \(\mathrm{RP}_{7}\) \\
\hline 0 & 8 & 16 & 24 & 32 & 40 & 48 & 56
\end{tabular}

Figure 13. Relative Priority Register
Each \(R P_{n}\) field is defined as follows.

\section*{Bits Meaning}

0:1 Reserved
2:7 Relative priority of priority level \(\mathbf{n}\) : Specifies the relative priority that corresponds to the priority corresponding to \(\mathrm{PPR}_{\text {PRI }}=\mathrm{n}\), where a value of 0 indicates the lowest relative priority and a value of \(0 b 111111\) indicates the highest relative priority.

\section*{Programming Note}

The hypervisor must ensure that the values of the \(R P_{n}\) fields increase monotonically for each \(n\) and are of different enough magnitudes to ensure that each priority level provides a meaningful difference in priority.

\subsection*{4.3.9 Software-use SPRs}

Software-use SPRs are 64-bit registers provided for use by software.
\begin{tabular}{|ll|}
\hline & SPRG0 \\
\hline SPRG1 \\
\hline SPRG2 \\
\hline 0 & SPRG3 \\
\hline
\end{tabular}

\section*{Figure 14. Software-use SPRs}

SPRG0, SPRG1, and SPRG2 are privileged registers. SPRG3 is a privileged register except that the contents may be copied to a GPR in Problem state when accessed using the mfspr instruction.

\section*{Programming Note}

Neither the contents of the SPRGs, nor accessing them using mtspr or mfspr, has a side effect on the operation of the thread. One or more of the registers is likely to be needed by non-hypervisor interrupt handler programs (e.g., as scratch registers and/or pointers to per thread save areas).
Operating systems must ensure that no sensitive data are left in SPRG3 when a problem state program is dispatched, and operating systems for secure systems must ensure that SPRG3 cannot be used to implement a "covert channel" between problem state programs. These requirements can be satisfied by clearing SPRG3 before passing control to a program that will run in problem state.

\section*{Version 3.0}

HSPRG0 and HSPRG1 are 64-bit registers provided for use by hypervisor programs.
\begin{tabular}{|ll|}
\hline \multicolumn{3}{|c|}{ HSPRG0 } \\
\hline & HSPRG1 \\
\hline 0 & 63
\end{tabular}

Figure 15. SPRs for use by hypervisor programs

\section*{Programming Note}

Neither the contents of the HSPRGs, nor accessing them using mtspr or mfspr, has a side effect on the operation of the thread. One or more of the registers is likely to be needed by hypervisor interrupt handler programs (e.g., as scratch registers and/or pointers to per thread save areas).

\subsection*{4.4 Fixed-Point Facility Instructions}

\subsection*{4.4.1 Fixed-Point Load and Store Caching Inhibited Instructions}

The storage accesses caused by the instructions described in this section are performed as though the specified storage location is Caching Inhibited and Guarded. The instructions can be executed only in hypervisor state. Software must ensure that the specified storage location is not in the caches. If the specified storage location is in a cache, the results are undefined.

The Fixed-Point Load and Store Caching Inhibited instructions must be executed only when \(\mathrm{MSR}_{\mathrm{DR}}=0\). The storage location specified by the instructions must not be in storage specified by the Hypervisor Real Mode Storage Control facility to be treated as
non-Guarded. If either of these conditions is violated, the result is a Data Storage interrupt.

\section*{Programming Note}

The instructions described in this section can be used to permit a control register on an I/O device to be accessed without permitting the corresponding storage location to be copied into the caches.

The Fixed-Point Load and Store Caching Inhibited instructions are fixed-point Storage Access instructions; see Section 3.3.1 of Book I.

\section*{Load Byte and Zero Caching Inhibited Indexed \\ X-form}

\section*{Ibzcix RT,RA,RB}
\begin{tabular}{|l|l|l|c|c|c|c|}
\hline 31 & \multicolumn{1}{|c|}{ RT } & RA & RB & \multicolumn{2}{|c|}{853} & \(/\) \\
0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else b
EA \leftarrow b + (RB)
RT \leftarrow }\mp@subsup{}{56}{0}0||MEM(EA, 1

```

Let the effective address (EA) be the sum (RAIO)+ (RB). The byte in storage addressed by EA is loaded into \(R T_{56: 63}\). \(\mathrm{RT}_{0: 55}\) are set to 0 .

The storage access caused by this instruction is performed as though the specified storage location is Caching Inhibited and Guarded.
This instruction is hypervisor privileged.
Special Registers Altered:
None

\section*{Load Word and Zero Caching Inhibited Indexed \\ \(X\)-form}

Iwzcix RT,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & & 789 & 1 \\
0 & & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
RT}\leftarrow\mp@subsup{}{}{32}0||MEM(EA, 4

```

Let the effective address (EA) be the sum (RAIO)+ (RB). The word in storage addressed by EA is loaded into \(R T_{32: 63} . \mathrm{RT}_{0: 31}\) are set to 0 .
The storage access caused by this instruction is performed as though the specified storage location is Caching Inhibited and Guarded.

This instruction is hypervisor privileged.

\section*{Special Registers Altered:}

None

\section*{Load Halfword and Zero Caching Inhibited Indexed \\ X-form}

Ihzcix RT,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{1}{|c|}{ RT } & RA & RB & \multicolumn{2}{|c|}{821} & \(/\) \\
0 & & & 11 & 16 & 21 & \\
\hline 1 \\
\hline
\end{tabular}
\[
\begin{aligned}
& \text { if } R A=0 \text { then } b \leftarrow 0 \\
& \text { else } \quad b \leftarrow \text { (RA) } \\
& E A \leftarrow b+(R B) \\
& R T \leftarrow{ }^{48} 0 \| \operatorname{MEM}(E A, 2)
\end{aligned}
\]

Let the effective address (EA) be the sum (RAIO)+ (RB). The halfword in storage addressed by \(E A\) is loaded into \(R T_{48: 63} . \mathrm{RT}_{0: 47}\) are set to 0 .

The storage access caused by this instruction is performed as though the specified storage location is Caching Inhibited and Guarded.

This instruction is hypervisor privileged.
Special Registers Altered:
None

\section*{Load Doubleword Caching Inhibited Indexed \\ X-form}

Idcix RT,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & RA & RB & & 885 & 1 \\
\hline 0 & 6 & 11 & 16 & 21 & & 31 \\
\hline
\end{tabular}
```

if RA = 0 then b \& 0
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
RT \leftarrowMEM(EA, 8)

```

Let the effective address (EA) be the sum (RAIO)+ (RB). The doubleword in storage addressed by EA is loaded into RT.

The storage access caused by this instruction is performed as though the specified storage location is Caching Inhibited and Guarded.
This instruction is hypervisor privileged.
Special Registers Altered: None

Store Byte Caching Inhibited Indexed X-form
stbcix
RS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 981 & \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
if \(\mathrm{RA}=0\) then \(\mathrm{b} \leftarrow 0\)
```

else b
EA \leftarrow b + (RB)
MEM(EA, 1) \leftarrow(RS) 56:63

```

Let the effective address (EA) be the sum (RAIO)+ (RB). (RS) 56:63 are stored into the byte in storage addressed by EA.

The storage access caused by this instruction is performed as though the specified storage location is Caching Inhibited and Guarded.

This instruction is hypervisor privileged.
Special Registers Altered:
None

\section*{Store Word Caching Inhibited Indexed \(X\)-form}
stwcix RS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & 917 & 1 \\
\hline 0 & 6 & 11 & 16 & 21 & 31 \\
\hline
\end{tabular}
```

if RA = 0 then b }\leftarrow
else b
EA \leftarrow b + (RB)
MEM(EA, 4) \leftarrow(RS) 32:63

```

Let the effective address (EA) be the sum (RAIO)+ (RB). (RS) \({ }_{32: 63}\) are stored into the word in storage addressed by EA.

The storage access caused by this instruction is performed as though the specified storage location is Caching Inhibited and Guarded.

This instruction is hypervisor privileged.
Special Registers Altered:
None

Store Halfword Caching Inhibited Indexed
X-form
sthcix RS,RA,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & RA & RB & & 949 & 1 \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}
```

if RA = 0 then b \leftarrow
else b
EA \leftarrow b + (RB)
MEM(EA, 2) \leftarrow (RS) 48:63

```

Let the effective address (EA) be the sum (RAIO)+ (RB). (RS) 48:63 are stored into the halfword in storage addressed by EA.

The storage access caused by this instruction is performed as though the specified storage location is Caching Inhibited and Guarded.

This instruction is hypervisor privileged.
Special Registers Altered:
None

\section*{Store Doubleword Caching Inhibited Indexed \\ X-form \\ stdcix RS,RA,RB \\ \begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \({ }_{6}\) RS & \multicolumn{1}{|c|}{ RA } & \({ }_{16}\) RB & & 1013 & 1 \\
0 & & 61 & & 31 \\
\hline
\end{tabular}}
```

if RA = 0 then b }\leftarrow
else b
EA}\leftarrow\textrm{b}+(\textrm{RB}
MEM(EA, 8) \leftarrow(RS)

```

Let the effective address (EA) be the sum (RAIO)+ (RB). (RS) is stored into the doubleword in storage addressed by EA.

The storage access caused by this instruction is performed as though the specified storage location is Caching Inhibited and Guarded.

This instruction is hypervisor privileged.
Special Registers Altered:
None

\subsection*{4.4.2 OR Instruction}
or \(R x, R x, R x\) can be used to set PPR \(\mathrm{PRI}_{\text {I }}\) (see Section 4.3.6) as shown in Figure 16. For all priorities except medium high, PPR \(_{\text {PRI }}\) remains unchanged if the privilege state of the thread executing the instruction is lower than the privilege indicated in the figure. For priority medium high, PPR \(_{\text {PRI }}\) is set to medium if the thread executing the instruction is in problem state and medium high priority is not allowed for problem state programs. (The encodings available to problem state programs, as well as encodings for additional shared resource hints not shown here, are described in Chapter 3 of Book II.)
\begin{tabular}{|c|c|l|c|}
\hline Rx & PPR \(\mathbf{P R I}^{\prime}\) & Priority & \begin{tabular}{l} 
Privi- \\
leged
\end{tabular} \\
\hline 31 & 001 & very low & no \\
\hline 1 & 010 & low & no \\
\hline 6 & 011 & medium low & no \\
\hline 2 & 100 & medium & no \\
\hline 5 & 101 & medium high & no/yes \({ }^{1}\) \\
\hline 3 & 110 & high & yes \\
\hline 7 & 111 & very high & hypv \\
\hline
\end{tabular}
\({ }^{1}\) This value is privileged unless the Problem State Priority Boost register allows the priority value 0b101 (See Section 4.3.7.)

Figure 16. Priority levels for or \(R x, R x, R x\)

\subsection*{4.4.3 Load Monitored Doubleword Instruction}

The Idmx instruction can only be executed in problem state. If an attempt is made to execute the instruction in privileged state, a Hypervisor Emulation Assistance interrupt will occur.

\section*{Programming Note}

In privileged state, if Idmx were not illegal it would behave as if it were Idx (because event-based exceptions and event-based branches do not occur in privileged state.) Such behavior could result in silent errors if privileged software used Idmx expecting it to behave as it does in problem state.

\subsection*{4.4.4 Transactional Memory Instructions}

Privileged software that makes the Transactional Memory Facility available to applications takes on the responsibility of managing the facility's resources and the application's transactional state during interrupt handling, service calls, task switches, and its own use of TM. In addition to the existing instructions like rfid and problem state TM instructions that play a role in this management, treclaim and trechkpt. may be used, as described below. See Section 3.2.2 for additional information about managing the TM facility and associated state transitions.

\section*{Transaction Reclaim \\ X-form}
treclaim. RA

```

CRO}\leftarrow0||\mp@subsup{MSR}{TS}{}||
if MSR
\#Transactional or Suspended
if RA = 0 then cause <- 0x00000001
else cause <- GPR(RA)56:63 || 0x000001
if TEXASR
Discard transactional footprint
TMRecordFailure(cause)
Revert checkpointed registers to pre-transac-
tional values
Discard all resources related to current
transaction
MSR

```

The treclaim. instruction frees the transactional facility for use by a new transaction. It sets condition register field 0 to 0 II \(\mathrm{MSR}_{\text {TS }}\) II 0 . If the transactional facility is in the Transactional state or Suspended state, failure recording is performed as defined in Section 5.3.2 of Book II. If RA is 0 , the failure cause is set to \(0 x 00000001\), otherwise it is set to \(\operatorname{GPR}(\mathrm{RA})_{56: 63}\) II \(0 x 000001\). The checkpointed registers are reverted to their pre-transactional values, and all resources related to the current transaction are discarded, including the transactional footprint (if it wasn't already discarded for a pending failure).

The transaction state is set to Non-transactional.
If an attempt is made to execute treclaim. in Non-transactional state, a TM Bad Thing type Program interrupt will be generated.

This instruction is privileged.

\section*{Special Registers Altered:}

CROTEXASR TFIAR TS

\section*{Programming Note}

The treclaim. instruction can be used by an interrupt handler to deallocate the current thread's transactional resources in preparation for subsequent use of the facility by a new transaction. (An abort is not appropriate for this use, because (a) the interrupt handler is in Suspended state and an abort in Suspended state leaves the thread in Suspended state, and (b) an abort in Suspended state does not restore the checkpointed registers to their pre-transaction values.) After treclaim. is executed, when the interrupted program is next dispatched it should be resumed by first using trechkpt. to restore the pre-transactional register values into the checkpoint area. Failure handling for that program will occur when the program next attempts to execute an instruction in the Transactional state, which will cause the failure handler to be invoked because TDOOMED will be 1. (This will be immediate if the program was in the Transactional state when the interrupt occurred, or will be after tresume. is executed if the program was in the Suspended state when the interrupt occurred.)

\section*{Transaction Recheckpoint \\ X-form}
trechkpt.
\begin{tabular}{|l|l|l|l|l|ll|c|}
\hline \multicolumn{1}{|c|}{31} & \multicolumn{1}{|c|}{\(/ / /\)} & \multicolumn{1}{|c|}{\(/ / /\)} & \multicolumn{1}{c|}{\(/ / /\)} & & 1006 & 1 \\
0 & & 6 & & 11 & 16 & 21 & \\
\hline
\end{tabular}
\[
\mathrm{CRO} 0 \leftarrow 0\left\|\mathrm{MSR}_{\mathrm{TS}}\right\| 0
\]
\(\mathrm{MSR}_{\mathrm{TS}} \leftarrow 0 \mathrm{Ob} 01\)
TDOOMED \(\leftarrow 1\)
checkpoint area \(\leftarrow\) (checkpointed registers)
The trechkpt. instruction copies the current (pre-transactional, saved and restored by the operating system) register state to the checkpoint area. It sets condition register field 0 to \(0\left\|\mathrm{MSR}_{\text {TS }}\right\| 0\). The current values of the checkpointed registers are loaded into the checkpoint area. TDOOMED is set to Ob1.

The transaction state is set to Suspended.
If an attempt is made to execute this instruction in Transactional or Suspended state or when TEXAS\(R_{\text {FS }}=0\), a TM Bad Thing type Program interrupt will be generated.

This instruction is privileged.

\section*{Special Registers Altered:}

CRO TS

\subsection*{4.4.5 Move To/From System Register Instructions}

The Move To Special Purpose Register and Move From Special Purpose Register instructions are described in Book I, but only at the level available to an application programmer. For example, no mention is made there of registers that can be accessed only in privileged state. The descriptions of these instructions given below extend the descriptions given in Book I, but do not list Special Purpose Registers that are implementa-tion-dependent. In the descriptions of these instructions given in below, the "defined" SPR numbers are the SPR numbers shown in the Figure 17 for the instruction and the implementation-specific SPR numbers that are implemented, and similarly for "defined" registers.
SPR 158, identified in Figure 17 as GSR, is a special SPR in that it retains no state and exists only to identify a performance optimization opportunity. mtspr specifying the GSR (Group Start Register) is used to identify the start of a sequence of mtspr instructions that may be optimized to have their SPR changes synchronized once as a group, rather than independently. The sequence is ended by any instruction other than a \(\boldsymbol{m t s p r}\) and also by an implicit redirection of instruction fetching, including those caused by interrupts, event-based branches, and transaction failure. This function may be useful when restoring a number of SPRs. If any of the mtspr instructions in the sequence
requires explicit context synchronization, a context synchronizing instruction must follow the sequence. See Chapter 11., "Synchronization Requirements for Context Alterations", on page 1127 for more details. mfspr specifying the GSR is a noop.

SPR numbers that are not shown in Figure 17 and are in the ranges shown below are reserved for implemen-tation-specific uses.

848-863
880-895
976-991
1008-1023
Implementation-specific registers must be privileged. SPR numbers for implementation-specific SPRs should be registered in advance with the Power ISA architects.

Figure 17. SPR encodings (Sheet 1 of 3)

I
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{decimal} & SPR \({ }^{1}\) & \multirow[b]{2}{*}{Register Name} & \multicolumn{2}{|r|}{Privileged} & \multirow[t]{2}{*}{Length (bits)} & \multicolumn{2}{|l|}{Extended Mnemonics*} \\
\hline & \(\mathbf{s p r}_{5: 9} \mathbf{s p r}_{0: 4}\) & & mtspr & mfspr & & mtspr & mfspr \\
\hline 1 & 0000000001 & XER & no & no & 64 & mtxer Rx & mfxer Rx \\
\hline 3 & 0000000011 & DSCR & no & no & 64 & mtudscr & mfudscr \\
\hline 8 & 0000001000 & LR & no & no & 64 & mtlr Rx & mflr Rx \\
\hline 9 & 0000001001 & CTR & no & no & 64 & mtctr Rx & mfttr Rx \\
\hline 13 & 0000001101 & AMR & no \({ }^{4}\) & no & 64 & mtuamr Rx & mfuamr Rx \\
\hline 17 & 0000010001 & DSCR & yes & yes & 64 & mtdscr Rx & mfdscr Rx \\
\hline 18 & 0000010010 & DSISR & yes & yes & 32 & mtdsisr Rx & mfdsisr Rx \\
\hline 19 & 0000010011 & DAR & yes & yes & 64 & mtdar Rx & mfdar Rx \\
\hline 22 & 0000010110 & DEC & yes & yes & 64 & mtdec Rx & mfdec Rx \\
\hline 26 & 0000011010 & SRR0 & yes & yes & 64 & mtsrr0 Rx & mfsrr0 Rx \\
\hline 27 & 0000011011 & SRR1 & yes & yes & 64 & mtsrr1 Rx & mfsrr1 Rx \\
\hline 28 & 0000011100 & CFAR & yes & yes & 64 & mtcfar Rx & mfcfar Rx \\
\hline 29 & 0000011101 & AMR & yes \({ }^{4}\) & yes & 64 & mtamr Rx & mfamr Rx \\
\hline 48 & 0000110000 & PIDR & yes & yes & 32 & mtpidr Rx & mfpidr Rx \\
\hline 61 & 0000111101 & IAMR & yes \({ }^{7}\) & yes & 64 & mtiamr Rx & mfiamr Rx \\
\hline 128 & 0010000000 & TFHAR & no & no & 64 & mttfhar Rx & mftfhar Rx \\
\hline 129 & 0010000001 & TFIAR & no & no & 64 & mttfiar Rx & mftfiar Rx \\
\hline 130 & 0010000010 & TEXASR & no & no & 64 & mttexasr Rx & mftexasr Rx \\
\hline 131 & 0010000011 & TEXASRU & no & no & 32 & mttexasru Rx & mftexasru Rx \\
\hline 136 & 0010001000 & CTRL & - & no & 32 & - & mfctrl Rx \\
\hline 152 & 0010011000 & CTRL & yes & - & 32 & mtctrl Rx & - \\
\hline 153 & 0010011001 & FSCR & yes & yes & 64 & mtfscr Rx & mffscr Rx \\
\hline 157 & 0010011101 & UAMOR & yes \(^{5}\) & yes & 64 & mtuamor Rx & mfuamor Rx \\
\hline 158 & 0010011110 & GSR & yes & yes & & mtgsr Rx & -12 \\
\hline 159 & 0010011111 & PSPB & yes & yes & 32 & mtpspb Rx & mfpspb Rx \\
\hline 176 & 0010110000 & DPDES & hypv \({ }^{2}\) & yes & 64 & mtdpdes Rx & mfdpdes Rx \\
\hline 180 & 0010110100 & DAWR0 & hypv \({ }^{2}\) & hypv \({ }^{2}\) & 64 & mtdawr0 Rx & mfdawr0 Rx \\
\hline 186 & 0010111010 & RPR & \(h^{\text {hpp }}{ }^{2}\) & \(h^{\text {ypp }}{ }^{2}\) & 64 & mtrpr Rx & mfrpr Rx \\
\hline 187 & 0010111011 & CIABR & \(h^{\text {hpv }}{ }^{2}\) & \(h^{\text {ypp }}{ }^{2}\) & 64 & mtciabr Rx & mfciabr Rx \\
\hline 188 & 0010111100 & DAWRX0 & \(h^{\text {hypv }}{ }^{2}\) & \(h^{\text {hypv }}{ }^{2}\) & 32 & mtdawrx0 Rx & mfdawrx0 Rx \\
\hline 190 & 0010111110 & HFSCR & hypv \({ }^{2}\) & \(h^{\text {ypv }}{ }^{2}\) & 64 & mthfscr Rx & mfhfscr Rx \\
\hline 256 & 0100000000 & VRSAVE & no & no & 32 & mtvrsave Rx & mfvrsave Rx \\
\hline 259 & 0100000011 & SPRG3 & - & no & 64 & - & mfusprg3 \\
\hline 268 & 0100001100 & TB & - & no & 64 & - & mftb Rx \({ }^{11}\) \\
\hline 269 & 0100001101 & TBU & - & no & 32 & - & mftbu Rx \({ }^{11}\) \\
\hline 272-275 & 01000 100xx & SPRG[n] n=0-3 & yes & yes & 64 & mtspgrn Rx & mfspgrn Rx \\
\hline 283 & 0100011011 & CIR & - & yes & 32 & - & mfcir Rx \\
\hline 284 & 0100011100 & TBL & \(h^{\prime \prime p v^{2}}\) & - & 32 & mttbl Rx & - \\
\hline 285 & 0100011101 & TBU & hypv \({ }^{2}\) & - & 32 & mttbu Rx & - \\
\hline 286 & 0100011110 & TBU40 & hypv & - & 64 & mttbu40 Rx & - \\
\hline 287 & 0100011111 & PVR & - & yes & 32 & - & mfpvr Rx \\
\hline 304 & 0100110000 & HSPRG0 & hypv \({ }^{2}\) & hypv \({ }^{2}\) & 64 & mthsprg0 Rx & mfhsprg0 Rx \\
\hline 305 & 0100110001 & HSPRG1 & \(h^{\text {ypv }}{ }^{2}\) & \(h^{\text {hypv }}{ }^{2}\) & 64 & mthsprg1 Rx & mfhsprg1 Rx \\
\hline 306 & 0100110010 & HDSISR & hypv \({ }^{2}\) & \(h^{\text {hypv }}{ }^{2}\) & 32 & mthdisr Rx & mfhdisr Rx \\
\hline 307 & 0100110011 & HDAR & \(h^{\text {hypv }}{ }^{2}\) & \(h^{\text {ypp }}{ }^{2}\) & 64 & mthdar Rx & mfhdar Rx \\
\hline 308 & 0100110100 & SPURR & \(h^{\text {hpv }}{ }^{2}\) & yes & 64 & mtspurr Rx & mfspurr Rx \\
\hline 309 & 0100110101 & PURR & hypv \({ }^{2}\) & yes & 64 & mtpurr Rx & mfpurr Rx \\
\hline 310 & 0100110110 & HDEC & \(h^{\text {hpp }}{ }^{2}\) & \(h^{\text {ypp }}{ }^{2}\) & 64 & mthdec Rx & mfhdec Rx \\
\hline 313 & 0100111001 & HRMOR & \(h^{\text {hypv }}{ }^{2}\) & \(h^{\text {hypv }}{ }^{2}\) & 64 & mthrmor Rx & mfhrmor Rx \\
\hline 314 & 0100111010 & HSRR0 & \(h^{\text {hpv }}{ }^{2}\) & \(h^{\text {hypv }}{ }^{2}\) & 64 & mthsrr0 Rx & mfhsrr0 Rx \\
\hline 315 & 0100111011 & HSRR1 & \(h^{\prime \prime p v}{ }^{2}\) & \(h^{\prime \prime p v}{ }^{2}\) & 64 & mthsrr1 Rx & mfhsrr1 Rx \\
\hline
\end{tabular}

Figure 17. SPR encodings (Sheet 2 of 3)

I
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multirow[b]{2}{*}{decimal} & SPR \({ }^{1}\) & \multirow[b]{2}{*}{Register Name} & \multicolumn{2}{|r|}{Privileged} & \multirow[t]{2}{*}{Length (bits)} & \multicolumn{2}{|l|}{Extended Mnemonics*} \\
\hline & \(\mathbf{s p r}_{5: 9} \mathbf{s p r}_{0: 4}\) & & mtspr & mfspr & & mtspr & mfspr \\
\hline 318 & 0100111110 & LPCR & hypv \({ }^{2}\) & hypv \({ }^{2}\) & 64 & mtlpcr Rx & mflper Rx \\
\hline 319 & 0100111111 & LPIDR & \(h^{\text {hypv }}{ }^{2}\) & \(h^{\text {ypp }}{ }^{2}\) & 32 & mtlpidr Rx & mflpidr Rx \\
\hline 336 & 0101010000 & HMER & \(h^{\text {yppv}}{ }^{2,3}\) & \(h^{\text {hypv }}{ }^{2}\) & 64 & mthmer Rx & mfhmer Rx \\
\hline 337 & 0101010001 & HMEER & hypv \({ }^{2}\) & \(h^{\text {ypp }}{ }^{2}\) & 64 & mthmeer Rx & mfhmeer Rx \\
\hline 338 & 0101010010 & PCR & hypv \({ }^{2}\) & \(h^{\text {hyp }}{ }^{2}\) & 64 & mtpcr Rx & mfpcr Rx \\
\hline 339 & 0101010011 & HEIR & hypv \({ }^{2}\) & \(h^{\text {hyp }}{ }^{2}\) & 32 & mtheir Rx & mfheir Rx \\
\hline 349 & 0101011101 & AMOR & \(h^{\prime \prime p v}{ }^{2}\) & \(h^{\text {ypv }}{ }^{2}\) & 64 & mtamor Rx & mfamor Rx \\
\hline 446 & 0110111110 & TIR & - & yes & 64 & - & mftir Rx \\
\hline 464 & 0111010000 & PTCR & \(h^{\prime} \mathrm{ypv}^{2}\) & \(h^{\text {ypv }}{ }^{2}\) & 64 & mtptcr Rx & mfptcr Rx \\
\hline 768 & 1100000000 & SIER & - & no \({ }^{6}\) & 64 & - & mfusier Rx mfsier Rx \\
\hline 769 & 1100000001 & MMCR2 & no \({ }^{6}\) & no \({ }^{6}\) & 64 & mtummcr2 Rx mtmmcr2 Rx & mfummcr2 Rx mfmmcr2 Rx \\
\hline 770 & 1100000010 & MMCRA & \(n 0^{6}\) & no \({ }^{6}\) & 64 & mtummcra Rx & mfummcra Rx mfmmcra Rx \\
\hline 771 & 1100000011 & PMC1 & no \({ }^{6}\) & no \({ }^{6}\) & 32 & mtupmc1 Rx & mfupmc1 Rx mfpmc1 Rx \\
\hline 772 & 1100000100 & PMC2 & no \({ }^{6}\) & no \({ }^{6}\) & 32 & mtupmc2 Rx & mfupmc2 Rx mfpmc2 Rx \\
\hline 773 & 1100000101 & PMC3 & \(n 0^{6}\) & \(n 0^{6}\) & 32 & mtupmc3 Rx & mfupmc3 Rx mfpmc3 Rx \\
\hline 774 & 1100000110 & PMC4 & \(n 0^{6}\) & no \({ }^{6}\) & 32 & mtupmc4 Rx & \begin{tabular}{l}
mfupmc4 Rx \\
mfpme4 Rx
\end{tabular} \\
\hline 775 & 1100000111 & PMC5 & no \({ }^{6}\) & no \({ }^{6}\) & 32 & mtupmc5 Rx & mfupmc5 Rx mfpmc5 Rx \\
\hline 776 & 1100001000 & PMC6 & no \({ }^{6}\) & no \({ }^{6}\) & 32 & mtupmc6 Rx & mfupmc6 Rx mfpmc6 Rx \\
\hline 779 & 1100001011 & MMCRO & \(n 0^{6}\) & \(n 0^{6}\) & 64 & mtummcr0 Rx & mfummcr0 Rx mfmmcr0 Rx \\
\hline 780 & 1100001100 & SIAR & - & no \({ }^{6}\) & 64 & - & mfusiar Rx mfsiar Rx \\
\hline 781 & 1100001101 & SDAR & - & no \({ }^{6}\) & 64 & - & \begin{tabular}{l}
mfusdar Rx \\
mfsdar Rx
\end{tabular} \\
\hline 782 & 1100001110 & MMCR1 & - & \(n 0^{6}\) & 64 & - & mfummer1 Rx mfmmer1 Rx \\
\hline 784 & 1100010000 & SIER & yes & yes & 64 & mtsier Rx & 13 \\
\hline 785 & 1100010001 & MMCR2 & yes & yes & 64 & 13 & 13 \\
\hline 786 & 1100010010 & MMCRA & yes & yes & 64 & mtmmcra Rx & 13 \\
\hline 787 & 1100010011 & PMC1 & yes & yes & 32 & mtpmc1 Rx & 13 \\
\hline 788 & 1100010100 & PMC2 & yes & yes & 32 & mtpmc2 Rx & 13 \\
\hline 789 & 1100010101 & PMC3 & yes & yes & 32 & mtpmc3 Rx & 13 \\
\hline 790 & 1100010110 & PMC4 & yes & yes & 32 & mtpmc4 Rx & 13 \\
\hline 791 & 1100010111 & PMC5 & yes & yes & 32 & mtpmc5 Rx & 13 \\
\hline 792 & 1100011000 & PMC6 & yes & yes & 32 & mtpmc6 Rx & 13 \\
\hline 795 & 1100011011 & MMCR0 & yes & yes & 64 & mtmmcr0 Rx & 13 \\
\hline 796 & 1100011100 & SIAR & yes & yes & 64 & mtsiar Rx & 13 \\
\hline 797 & 1100011101 & SDAR & yes & yes & 64 & mtsdar Rx & 13 \\
\hline 798 & 1100011110 & MMCR1 & yes & yes & 64 & mtmmcr1 Rx & 13 \\
\hline 800 & 1100100000 & BESCRS & no & no & 64 & mtbescrs Rx & mfbescrs Rx \\
\hline 801 & 1100100001 & BESCRSU & no & no & 32 & mtbescrsu Rx & mfbescrsu Rx \\
\hline 802 & 1100100010 & BESCRR & no & no & 64 & mtbescrr Rx & mfbescrr Rx \\
\hline 803 & 1100100011 & BESCRRU & no & no & 32 & mtbescrru Rx & mfbescrru Rx \\
\hline 804 & 1100100100 & EBBHR & no & no & 64 & mtebbhr Rx & mfebbhr Rx \\
\hline 805 & 1100100101 & EBBRR & no & no & 64 & mtebbrr Rx & mfebbrr Rx \\
\hline 806 & 1100100110 & BESCR & no & no & 64 & mtbescr Rx & mfbescr Rx \\
\hline 808 & 1100101000 & reserved \({ }^{8}\) & no & no & na & - & - \\
\hline
\end{tabular}

Figure 17. SPR encodings (Sheet 3 of 3)
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multirow[t]{2}{*}{decimal} & SPR \({ }^{1}\) & \multirow[b]{2}{*}{Register Name} & \multicolumn{2}{|r|}{Privileged} & \multirow[t]{2}{*}{Length (bits)} & \multicolumn{2}{|l|}{Extended Mnemonics*} \\
\hline & \(\mathbf{s p r}_{5: 9} \mathbf{s p r}_{0: 4}\) & & mtspr & mfspr & & mtspr & mfspr \\
\hline 809 & 1100101001 & reserved \({ }^{8}\) & no & no & na & - & - \\
\hline 810 & 1100101010 & reserved \({ }^{8}\) & no & no & na & - & - \\
\hline 811 & 1100101011 & reserved \({ }^{8}\) & no & no & na & - & - \\
\hline 813 & 1100101011 & LMRR & no & no & 64 & mtlmrr Rx & mflmrr Rx \\
\hline 814 & 1100101101 & LMSER & no & no & 64 & mtlmser Rx & mflmser Rx \\
\hline 815 & 1100101110 & TAR & no & no & 64 & mttar Rx & mftar Rx \\
\hline 816 & 1100110000 & ASDR & \(h^{\text {ypv }}{ }^{2}\) & \(h^{\text {ypv }}{ }^{2}\) & 64 & mtasdr Rx & mfasdr Rx \\
\hline 823 & 1100110111 & PSSCR & yes & yes & 64 & mtpsscr Rx & mfpsscr Rx \\
\hline 848 & 1101010000 & IC & \(h^{\text {hypv }}{ }^{2}\) & yes & 64 & mtic Rx & mfic Rx \\
\hline 849 & 1101010001 & VTB & hypv2 & yes & 64 & mtvtb Rx & mfvtb Rx \\
\hline 855 & 1101010111 & PSSCR & \(h^{\text {hpv }}{ }^{3}\) & \(\mathrm{hypv}^{3}\) & 64 & mthpsscr Rx & mfhpsscr \\
\hline 896 & 1110000000 & PPR & no & no & 64 & mtppr Rx & mfppr Rx \\
\hline 898 & 1110000010 & PPR32 & no & no & 32 & mtppr32 Rx & mfppr32 Rx \\
\hline 1023 & 1111111111 & PIR & - & yes & 32 & - & mfpir Rx \\
\hline
\end{tabular}
- This register is not defined for this instruction.

Note that the order of the two 5-bit halves of the SPR number is reversed.
2 This register is a hypervisor resource, and can be accessed by this instruction only in hypervisor state (see Chapter 2).
3 This register cannot be directly written. Instead, bits in the register corresponding to 0 bits in (RS) can be cleared using mtspr SPR,RS.
4 The value specified in register RS may be masked by the contents of the [U]AMOR before being placed into the AMR; see the mtspr instruction description.
5 The value specified in register RS may be ANDed with the contents of the AMOR before being placed into the UAMOR; see the mtspr instruction description.
6 MMCRO \(_{\text {PMCC }}\) controls the availability of this SPR, and its contents depend on the privilege state in which it is accessed. See Section 9.4.4 for details.
7 The value specified in Register RS may be masked by the contents of the AMOR before being placed into the IAMR; see the mtspr instruction description.
8 Accesses to these SPRs are noops; see Section 1.3.3, "Reserved Fields, Reserved Values, and Reserved SPRs" in Book I.
9 The length of the GSR is undefined. An access to this SPR affects synchronization of subsequent mtspr instructions. See the introductory text in this section for more details
\({ }^{10}\) SPR numbers 777-778, 783, 793-794, and 799 are reserved for the Performance Monitor. All other SPR numbers that are not shown above and are not implementation-specific are reserved.
11 The mftb instruction is Phased-Out. Assemblers targeting Version 2.03 or later of the architecture should generate an \(\boldsymbol{m f s p r}\) instruction for the \(\boldsymbol{m f t b}\) and \(\boldsymbol{m f t b u}\) extended mnemonics; see the corresponding Assembler Note in the mftb instruction description (see Section 6.1 of Book II).
\({ }^{12} \boldsymbol{m f s p r}\) specifying the GSR has no meaningful use. It is treated as a noop. As a result, no extended mnemonic is assigned for it.
\({ }^{13}\) No extended mnemonic is provided because previous versions of the architecture defined the obvious extended mnemonic as resolving to the non-privileged SPR number, and because there is no software benefit in using the privileged SPR number, rather than the non-privileged SPR number, for this function.
*This figure also defines extended mnemonics for the mtspr and mfspr instructions, including the Special Purpose

The mtspr and mfspr instructions specify an SPR as a numeric operand; extended mnemonics are provided that represent the SPR in the mnemonic rather than requiring it to be coded as an operand. Similar extended mnemonics are provided for the Move From Time Base instruction, which specifies the portion of the Time Base as a numeric operand.

Note: mftb serves as both a basic and an extended mnemonic. The Assembler will recognize an mftb mnemonic with two operands as the basic form, and an mftb mnemonic with one operand as the extended form. In the extended form the TBR operand is omitted and assumed to be 268 (the value that corresponds to TB)

\section*{Move To Special Purpose Register} XFX-form

The SPR field denotes a Special Purpose Register, encoded as shown in Figure 17. If the SPR field contains the value 158, the instruction indicates the start of a sequence of \(\boldsymbol{m t s p r}\) instructions that may be synchronized as a group. See the introductory material in this section for more information. If the SPR field contains a value from 808 through 811 , the instruction specifies a reserved SPR, and is treated as a no-op; see Section 1.3.3, "Reserved Fields, Reserved Values, and Reserved SPRs" in Book I. Otherwise, the contents of register RS are placed into the designated Special Purpose Register, except as described in the next four paragraphs. For Special Purpose Registers that are 32 bits long, the low-order 32 bits of RS are placed into the SPR.

When the designated SPR is the Authority Mask Register (AMR), (using SPR 13 or SPR 29), or the designated SPR is the Instruction Authority Mask Register (IAMR), and \(M S R_{H V ~ P R}=0 b 00\), the contents of bit positions of register RS corresponding to 1 bits in the Authority Mask Override Register (AMOR) are placed into the corresponding bits of the AMR or IAMR, respectively; the other AMR or IAMR bits are not modified.

When the designated SPR is the AMR, using SPR 13, and \(M_{2 S R}=1\), the contents of bit positions of register RS corresponding to 1 bits in the User Authority Mask Override Register (UAMOR) are placed into the corresponding bits of the AMR; the other AMR bits are not modified.
When the designated SPR is the UAMOR and \(M_{S R}{ }_{H V P R}=0 b 00\), the contents of register RS are ANDed with the contents of the AMOR and the result is placed into the UAMOR.
When the designated SPR is the Hypervisor Maintenance Exception Register (HMER), the contents of register RS are ANDed with the contents of the HMER and the result is placed into the HMER.

For this instruction, SPRs TBL and TBU are treated as separate 32-bit registers; setting one leaves the other unaltered.
\(\mathrm{spr}_{0}=1\) if and only if writing the register is privileged. Execution of this instruction specifying an SPR number with \(\mathrm{spr}_{0}=1\) causes a Privileged Instruction type Program interrupt when \(\mathrm{MSR}_{\mathrm{PR}}=1\) and, if the SPR is a hypervisor resource (see Figure 17) when \(M^{\prime} R_{H V ~ P R}=0 b 00\), causes a Privileged Instruction type of Program interrupt if \(\mathrm{LPCR}_{\text {EVIRT }}=0\) and a Hypervisor Emulation Assistance interrupt if \(\mathrm{LPCR}_{\mathrm{EVIRT}}=1\). .

Execution of this instruction specifying an SPR number that is not defined for the implementation causes one of the following.
- if \(\mathrm{spr}_{0}=0\) :
- if \(\mathrm{MSR}_{\mathrm{PR}}=1\) : Hypervisor Emulation Assistance interrupt
- if \(\mathrm{MSR}_{\mathrm{PR}}=0\) : Hypervisor Emulation Assistance interrupt for SPR 0 and no operation (i.e., the instruction is treated as a no-op) when LPCR \(_{\text {EVIRT }}=0\) or Hypervisor Emulation Assistance interrupt when LPCR \(_{\text {EVIRT }}=1\) for all other SPRs
- if \(\mathrm{spr}_{0}=1\) :
- if \(M S R_{P R}=1\) : Privileged Instruction type Program interrupt
- if \(M S R_{P R}=0\) : no operation (i.e., the instruction is treated as a no-op) when LPCR \({ }_{\text {EVIRT }}=0\) and Hypervisor Emulation Assistance interrupt when LPCR \(_{\text {EVIRT }}=1\)
If an attempt is made to execute mtspr specifying a Transactional Memory SPR in other than Non-transactional state, with the exception of TFAR in suspended state, a TM Bad Thing type Program interrupt is generated.

\section*{Special Registers Altered:}

See Figure 17

\section*{Programming Note \\ For a discussion of software synchronization requirements when altering certain Special Purpose Registers, see Chapter 11. "Synchronization Requirements for Context Alterations" on page 1127.}

\section*{Programming Note}

An attempt to execute an mtspr or mfspr instruction with SPR=0 or an attempt to execute an \(\boldsymbol{m f s p r}\) instruction with SPR=4, 5 , or 6 results in a Hypervisor Emulation Assistance interrupt to enable efficient emulation of \(\boldsymbol{m t} / f s p r\) specifying the corresponding SPRs as defined in the POWER Architecture.

\section*{Move From Special Purpose Register \\ XFX-form}


The SPR field denotes a Special Purpose Register, encoded as shown in Figure 17. If the designated Special Purpose Register is the TFIAR and TFIAR indicates the failure was recorded in a state more privileged than the current state, register RT is set to zero. If the SPR field contains 158 , the instruction specifies the GSR, and is treated as a noop. If the SPR field contains a value from 808 through 811, the instruction specifies a reserved SPR, and is treated as a no-op; see Section 1.3.3, "Reserved Fields, Reserved Values, and Reserved SPRs" in Book I. Otherwise, the contents of the designated Special Purpose Register are placed into register RT. For Special Purpose Registers that are 32 bits long, the low-order 32 bits of RT receive the contents of the Special Purpose Register and the high-order 32 bits of RT are set to zero.

\section*{Programming Note}

Note that when a problem state transaction's failure is recorded in hypervisor state and there is a subsequent need for a context switch in privileged, non-hypervisor state, an attempt to save TFIAR will result in zeros being saved. This is harmless because if the original application ever tries to read the TFIAR, it would read zeros anyway, since the failure took place in hypervisor state.
\(\mathrm{spr}_{0}=1\) if and only if reading the register is privileged. Execution of this instruction specifying an SPR number with \(\mathrm{spr}_{0}=1\) causes a Privileged Instruction type Program interrupt when \(M_{P R}=1\) and, if the SPR is a hypervisor resource (see Figure 17) when \(M_{\text {MSR }}^{H V}\) PR \(=0 b 00\), causes a Privileged Instruction type of Program interrupt when LPCR EVIRT \(=0\) and a Hypervisor Emulation Assistance interrupt when LPCR \(_{E-}\) VIRT=1.

\section*{Version 3.0}

Execution of this instruction specifying an SPR number that is not defined for the implementation causes one of the following.
- if \(\mathrm{spr}_{0}=0\) :
- if \(\mathrm{MSR}_{\mathrm{PR}}=1\) : Hypervisor Emulation Assistance interrupt
- if \(\mathrm{MSR}_{\mathrm{PR}}=0\) : Hypervisor Emulation Assistance interrupt for SPRs 0, 4, 5, and 6 and no operation (i.e., the instruction is treated as a no-op) when LPCR \(_{\text {EVIRT }}=0\) and Hypervisor Emulation Assistance interrupt when LPCR EVIRT \(=1\) for all other SPRs
- if \(\mathrm{spr}_{0}=1\) :
- if \(M S R_{P R}=1\) : Privileged Instruction type Program interrupt
- if \(\mathrm{MSR}_{\mathrm{PR}}=0\) : no operation (i.e., the instruction is treated as a no-op) when \(\operatorname{LPCR}_{\text {EVIRT }}=0\) and Hypervisor Emulation Assistance interrupt when LPCR \(_{\text {EVIRT }}=1\)

\section*{Special Registers Altered:}

None
Note
See the Notes that appear with mtspr.

\section*{Move To Machine State Register X-form}
mtmsr RS,L
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 31 & RS & I/I & L & I/I & & 146 & \(/\) \\
\hline 0 & & & 6 & 11 & 15 & 16 & \\
\hline
\end{tabular}
```

if $\mathrm{L}=0$ then
$\mathrm{MSR}_{48} \leftarrow(\mathrm{RS})_{48} \mid(\mathrm{RS})_{49}$
$\mathrm{MSR}_{58} \leftarrow(\mathrm{RS})_{58} \quad(\mathrm{RS})_{49}$
$\mathrm{MSR}_{59} \leftarrow(\mathrm{RS})_{59} \mid(\mathrm{RS})_{49}$
$\mathrm{MSR}_{32: 47} 49: 5052: 57$ 60:62 $\leftarrow(\mathrm{RS})_{32: 47} 49: 50 \quad 52: 57$ 60:62
else
$\mathrm{MSR}_{48} 62 \leftarrow(\mathrm{RS})_{48} 62$

```

The MSR is set based on the contents of register RS and of the \(L\) field.
\(\mathrm{L}=0\) :
The result of ORing bits 48 and 49 of register RS is placed into \(\mathrm{MSR}_{48}\). The result of ORing bits 58 and 49 of register RS is placed into \(\mathrm{MSR}_{58}\). The result of ORing bits 59 and 49 of register RS is placed into MSR \(_{59}\). Bits 32:47, 49:50, 52:57, and 60:62 of register RS are placed into the corresponding bits of the MSR.
\(\mathrm{L}=1\) :
Bits 48 and 62 of register RS are placed into the corresponding bits of the MSR. The remaining bits of the MSR are unchanged.
This instruction is privileged.
If \(L=0\) this instruction is context synchronizing. If \(L=1\) this instruction is execution synchronizing; in addition, the alterations of the EE and RI bits take effect as soon as the instruction completes.

\section*{Special Registers Altered: \\ MSR}

Except in the mtmsr instruction description in this section, references to "mtmsr" in this document imply either \(L\) value unless otherwise stated or obvious from context (e.g., a reference to an mtmsr instruction that modifies an MSR bit other than the EE or RI bit implies \(\mathrm{L}=0\) ).

\section*{Programming Note}

If this instruction sets \(M S R_{P R}\) to 1 , it also sets \(\mathrm{MSR}_{\mathrm{EE}}, \mathrm{MSR}_{\mathrm{IR}}\), and \(\mathrm{MSR}_{\mathrm{DR}}\) to 1 .

This instruction does not alter \(M_{\text {MSR }}^{\text {ME }}\) or MSR \(_{\text {LE }}\). (This instruction does not alter \(\mathrm{MSR}_{\mathrm{HV}}\) because it does not alter any of the high-order 32 bits of the MSR.)
If the only MSR bits to be altered are \(M_{\text {SR }}^{\text {EE RI }}\), to obtain the best performance \(\mathrm{L}=1\) should be used.

\section*{Programming Note}

If \(\mathrm{MSR}_{\mathrm{EE}}=0\) and an External, Decrementer, or Performance Monitor exception is pending, executing an mtmsrd instruction that sets \(\mathrm{MSR}_{\mathrm{EE}}\) to 1 will cause the interrupt to occur before the next instruction is executed, if no higher priority exception exists (see Section 6.9, "Interrupt Priorities" on page 1086). Similarly, if a Hypervisor Decrementer interrupt is pending, execution of the instruction by the hypervisor causes a Hypervisor Decrementer interrupt to occur if HDICE=1.

For a discussion of software synchronization requirements when altering certain MSR bits, see Chapter 11.

\section*{Programming Note}
mtmsr serves as both a basic and an extended mnemonic. The Assembler will recognize an mtmsr mnemonic with two operands as the basic form, and an mtmsr mnemonic with one operand as the extended form. In the extended form the L operand is omitted and assumed to be 0 .

\section*{Programming Note}

There is no need for an analogous version of the mfmsr instruction, because the existing instruction copies the entire contents of the MSR to the selected GPR.

\section*{Move To Machine State Register Doubleword}

X-form mtmsrd RS,L
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline 31 & RS & & I/I & L & & I/I & & 178 \\
0 & & 6 & & 11 & 15 & 16 & & 21 \\
\hline
\end{tabular}
```

if $\mathrm{L}=0$ then
if ( MSR $_{29: 31} ᄀ=0 \mathrm{~b} 010 \mid \mathrm{RS}_{29: 31} \mathrm{I}=0 \mathrm{~b} 000$ ) then
$\mathrm{MSR}_{29: 31} \leftarrow \mathrm{RS}_{29: 31}$
$\mathrm{MSR}_{48} \leftarrow(\mathrm{RS})_{48} \mid(\mathrm{RS})_{49}$
$\mathrm{MSR}_{58} \leftarrow(\mathrm{RS})_{58} \quad(\mathrm{RS})_{49}$
$\mathrm{MSR}_{59} \leftarrow(\mathrm{RS})_{59} \mid(\mathrm{RS})_{49}$
MSR $_{0}: 2$ 4:28 32:47 49:50 52:57 60:62
$\leftarrow(\text { RS })_{0: 2} 4$ 6:28 32:47 49:50 52:57 60:62
else
MSR $_{48} 62 \leftarrow(\mathrm{RS}) 4862$

```

The MSR is set based on the contents of register RS and of the \(L\) field.
\(\mathrm{L}=0\) :
If bits 29 through 31 of the MSR are not equal to 0b010 or bits 29 through 31 of RS are not equal to 0b000, then the value of bits 29 through 31 of RS is placed into bits 29 through 31 of the MSR.The result of ORing bits 48 and 49 of register RS is placed into \(\mathrm{MSR}_{48}\). The result of ORing bits 58 and 49 of register RS is placed into \(\mathrm{MSR}_{58}\). The result of ORing bits 59 and 49 of register RS is placed into \(\mathrm{MSR}_{59}\). Bits \(0: 2,4: 28,32: 47,49: 50,52: 57\), and \(60: 62\) of register RS are placed into the corresponding bits of the MSR.
\(\mathrm{L}=1\) :
Bits 48 and 62 of register RS are placed into the corresponding bits of the MSR. The remaining bits of the MSR are unchanged.

If the instruction attempts to cause an illegal transaction state transition (see Table 3, "Transaction state transitions that can be requested by rfebb, rfid, rfscv, hrfid, and mtmsrd.," on page 947), or when TM is disabled by the PCR, a transition to Problem state with an active transaction, a TM Bad Thing type Program interrupt is generated (unless a higher-priority exception is pending). If this interrupt is generated, the value placed into SRRO by the interrupt processing mechanism (see Section 6.4.3) is the address of the mtmsrd instruction.
This instruction is privileged.
If \(L=0\) this instruction is context synchronizing. If \(L=1\) this instruction is execution synchronizing; in addition, the alterations of the EE and RI bits take effect as soon as the instruction completes.

\section*{Special Registers Altered:}

MSR

Except in the mtmsrd instruction description in this section, references to "mtmsrd" in this document imply either \(L\) value unless otherwise stated or obvious from context (e.g., a reference to an mtmsrd instruction that modifies an MSR bit other than the EE or RI bit implies \(\mathrm{L}=0\) ).

\section*{Programming Note}

If this instruction sets \(M_{\text {M }} R_{P R}\) to 1 , it also sets \(M S R_{E E}, M S R_{I R}\), and \(M S R_{D R}\) to 1 .

This instruction does not alter \(\mathrm{MSR}_{\text {LE }}, \mathrm{MSR}_{\text {ME }}\) or \(\mathrm{MSR}_{\mathrm{HV}}\).
If the only MSR bits to be altered are MSR EE RI, , to obtain the best performance \(L=1\) should be used.

\section*{Programming Note}

If \(\mathrm{MSR}_{\mathrm{EE}}=0\) and an External, Decrementer, or Performance Monitor exception is pending, executing an mtmsrd instruction that sets \(\mathrm{MSR}_{\text {EE }}\) to 1 will cause the interrupt to occur before the next instruction is executed, if no higher priority exception exists (see Section 6.9, "Interrupt Priorities" on page 1086). Similarly, if a Hypervisor Decrementer interrupt is pending, execution of the instruction by the hypervisor causes a Hypervisor Decrementer interrupt to occur if HDICE=1.

For a discussion of software synchronization requirements when altering certain MSR bits, see Chapter 11.

\section*{Programming Note}
mtmsrd serves as both a basic and an extended mnemonic. The Assembler will recognize an mtmsrd mnemonic with two operands as the basic form, and an mtmsrd mnemonic with one operand as the extended form. In the extended form the L operand is omitted and assumed to be 0 .

\section*{Move From Machine State Register}

X-form
mfmsr RT

\(\mathrm{RT} \leftarrow \mathrm{MSR}\)
The contents of the MSR are placed into register RT.
This instruction is privileged.
Special Registers Altered:
None

Version 3.0

\title{
Chapter 5. Storage Control
}

\subsection*{5.1 Overview}

A program references storage using the effective address computed by the hardware when it executes a Load, Store, Branch, or Cache Management instruction, or when it fetches the next sequential instruction. The effective address is translated to a real address according to procedures described in Section 5.7.3, in Section 5.7.7 and in the following sections. The real address is what is presented to the storage subsystem.

For a complete discussion of storage addressing and effective address calculation, see Section 1.11 of Book I.

\subsection*{5.2 Storage Exceptions}

A storage exception results when the sequential execution model requires that a storage access be performed but the access is not permitted (e.g., is not permitted by the storage protection mechanism), the access cannot be performed because the effective address cannot be translated to a real address, or the access matches some tracking mechanism criteria (e.g., Data Address Watchpoint).

In certain cases a storage exception may result in the "restart" of (re-execution of at least part of) a Load or Store instruction. See Section 2.2 of Book II, and Section 6.6 in this Book.

\subsection*{5.3 Instruction Fetch}

Instructions are fetched under control of \(\mathrm{MSR}_{\mathrm{IR}}\).
\(\mathrm{MSR}_{\text {IR }}=\mathbf{0}\)
The effective address of the instruction is interpreted as described in Section 5.7.3.

MSR \(_{\text {IR }}=1\)
The effective address of the instruction is translated by the Address Translation mechanism described beginning in Section 5.7.7.

\subsection*{5.3.1 Implicit Branch}

Explicitly altering certain MSR bits (using mtmsr[d]), or explicitly altering SLB entries, Page Table Entries, or certain System Registers (including the HRMOR, and possibly other implementation-dependent registers), may have the side effect of changing the addresses, effective or real, from which the current instruction stream is being fetched. This side effect is called an implicit branch. For example, an mtmsrd instruction that changes the value of \(M_{\text {SF }}\) may change the effective addresses from which the current instruction stream is being fetched. The MSR bits and System Registers (excluding implementation-dependent registers) for which alteration can cause an implicit branch are indicated as such in Chapter 11. "Synchronization Requirements for Context Alterations" on page 1127. Implicit branches are not supported by the Power ISA. If an implicit branch occurs, the results are boundedly undefined.

\subsection*{5.3.2 Address Wrapping Combined with Changing MSR Bit SF}

If the current instruction is at effective address \(2^{32}-4\) and is an mtmsrd instruction that changes the contents of \(\mathrm{MSR}_{\text {SF }}\) the effective address of the next sequential instruction is undefined.

\section*{Programming Note}

In the case described in the preceding paragraph, if an interrupt occurs before the next sequential instruction is executed, the contents of SRRO, or HSRRO, as appropriate to the interrupt, are undefined.

\subsection*{5.4 Data Access}

Data accesses are controlled by \(\mathrm{MSR}_{\text {DR }}\).
\(M_{\text {DR }}=0\)
The effective address of the data is interpreted as described in Section 5.7.3.
\(\mathrm{MSR}_{\mathrm{DR}}=1\)
The effective address of the data is translated by the Address Translation mechanism described in Section 5.7.7.

\subsection*{5.5 Performing Operations Out-of-Order}

An operation is said to be performed "in-order" if, at the time that it is performed, it is known to be required by the sequential execution model. An operation is said to be performed "out-of-order" if, at the time that it is performed, it is not known to be required by the sequential execution model.

Operations are performed out-of-order on the expectation that the results will be needed by an instruction that will be required by the sequential execution model. Whether the results are really needed is contingent on everything that might divert the control flow away from the instruction, such as Branch, Trap, System Call, and Return From Interrupt instructions, and interrupts, and on everything that might change the context in which the instruction is executed.

Typically, operations are performed out-of-order when resources are available that would otherwise be idle, so the operation incurs little or no cost. If subsequent events such as branches or interrupts indicate that the operation would not have been performed in the sequential execution model, any results of the operation are abandoned (except as described below).
In the remainder of this section, including its subsections, "Load instruction" includes the Cache Management and other instructions that are stated in the instruction descriptions to be "treated as a Load", and similarly for "Store instruction".

A data access that is performed out-of-order may correspond to an arbitrary Load or Store instruction (e.g., a Load or Store instruction that is not in the instruction stream being executed). Similarly, an instruction fetch that is performed out-of-order may be for an arbitrary instruction (e.g., the aligned word at an arbitrary location in instruction storage).

Most operations can be performed out-of-order, as long as the machine appears to follow the sequential execution model. Certain out-of-order operations are restricted, as follows.
- Stores

Stores are not performed out-of-order (even if the Store instructions that caused them were executed out-of-order).
- Accessing Guarded Storage

The restrictions for this case are given in Section 5.8.1.1.

The only permitted side effects of performing an operation out-of-order are the following.
- A Machine Check or Checkstop that could be caused by in-order execution may occur out-of-order.
- Reference and Change bits may be set as described in Section 5.7.13.
■ Non-Guarded storage locations that could be fetched into a cache by in-order fetching or execution of an arbitrary instruction may be fetched out-of-order into that cache.

\subsection*{5.6 Invalid Real Address}

A storage access (including an access that is performed out-of-order; see Section 5.5) may cause a Machine Check if the accessed storage location contains an uncorrectable error or does not exist.
In the case that the accessed storage location does not exist, the Checkstop state may be entered. See Section 6.5.2 on page 1064.

\section*{Programming Note}

In configurations supporting multiple partitions, hypervisor software must ensure that a storage access by a program in one partition will not cause a Checkstop or other system-wide event that could affect the integrity of other partitions (see Chapter 2). For example, such an event could occur if a real address placed in a Page Table Entry does not exist.

\subsection*{5.7 Storage Addressing}

\section*{Storage Control Overview}

I - Host real address space size is \(2^{m}\) bytes, \(m \leq 60\); see Note 1.
■ Guest real address space size is \(2^{m}\) bytes, \(\mathrm{m} \leq 60\); see Notes 1 and 2.
- Real page size is \(2^{12}\) bytes ( 4 KB ).
- Effective address space size is \(2^{64}\) bytes.
- An effective address is translated to a virtual address via the Segment Lookaside Buffer (SLB).
- Virtual address space size is \(2^{n}\) bytes, | \(65 \leq n \leq 78\); see Note 3 .
- \(\quad\) Segment size is \(2^{s}\) bytes, \(s=28\) or 40 .
- \(\quad 2^{n-40} \leq\) number of virtual segments \(\leq 2^{n-28}\); see Note 3.
- Virtual page size is \(2^{p}\) bytes, where \(12 \leq p\), and \(2^{p}\) is no larger than either the size of the biggest segment or the real address space; a size of \(4 \mathrm{~KB}, 64 \mathrm{~KB}, 2 \mathrm{MB}\) for Radix Tree translation, and an implementation-dependent number of other sizes are supported; see Note 4. For HPT translation, the Page Table specifies the virtual page size and the SLB or implied segment descriptor specifies the base virtual page size, which is the smallest virtual page size that the segment can contain. The base virtual page size is \(2^{\text {b }}\) bytes. For Radix Tree translation, the virtual page size is determined by the location of the Page Table Entry in the Radix Tree.
- Segments contain pages of a single size, a mixture of 4 KB and 64 KB pages, or a mixture of page sizes that include implementa-tion-dependent page sizes.
- A virtual address is translated to a real address via the Page Table.

\section*{Notes:}
1. The value of \(m\) is implementation-dependent (subject to the maximum given above). When used to address storage or to represent a guest real address, the high-order \(60-\mathrm{m}\) bits of the " 60 -bit" real address must be zeros.
2. The hypervisor may assign a guest real address space size for each partition. Accesses to guest real storage outside this range but still mappable by the second level Radix Tree or HPT will cause an HDSI. Accesses to storage outside the mappable range will have boundedly undefined results.
3. The value of \(n\) is implementation-dependent (subject to the range given above). In references to 78 -bit virtual addresses elsewhere in this Book, the
high-order 78-n bits of the "78-bit" virtual address are assumed to be zeros.
4. The supported values of \(p\) for the larger virtual page sizes are implementation-dependent (subject to the limitations given above).

\section*{- Programming Note}

Note that without some of the reserved bits in the Radix PTE, the RPN field cannot address the full 60-bit real address space. Similarly without some of the reserved bits in the HPT PTE, the AVA field cannot address the full 60-bit real address space.

Note that without some of the reserved bits in the Radix PTE, the AVA field cannot resolve the full 78-bit virtual address.

\subsection*{5.7.1 32-Bit Mode}

The computation of the 64-bit effective address is independent of whether the thread is in 32-bit mode or 64-bit mode. In 32-bit mode ( \(\mathrm{MSR}_{\mathrm{SF}}=0\) ), the high-order 32 bits of the 64-bit effective address are treated as zeros for the purpose of addressing storage. This applies to both data accesses and instruction fetches. It applies independent of whether address translation is enabled or disabled. This truncation of the effective address is the only respect in which storage accesses in 32 -bit mode differ from those in 64 -bit mode.

\section*{Programming Note}

Treating the high-order 32 bits of the effective address as zeros effectively truncates the 64-bit effective address to a 32-bit effective address such as would have been generated on a 32-bit implementation of the Power ISA. Thus, for example, the ESID in 32-bit mode is the high-order four bits of this truncated effective address; the ESID thus lies in the range \(0-15\). When address translation is enabled, these four bits would select a Segment Register on a 32-bit implementation of the Power ISA. The SLB entries that translate these 16 ESIDs can be used to emulate these Segment Registers.

\subsection*{5.7.2 Virtualized Partition Memory (VPM) Mode}

VPM mode enables the hypervisor to reassign all or part of a partition's memory transparently so that the reassignment is not visible to the partition. When this is done, the partition's memory is said to be "virtualized." This mode is only available within Paravirtualized HPT translation mode. The other translation mode provides equivalent function by providing two levels of translation
with separate Page Tables for the operating system and the hypervisor. (See Section 5.7.7 for a more complete overview of the translation modes.) The VPM field in the LPCR enables VPM mode when address translation is enabled. VPM is always enabled when address translation is disabled.

If the thread is not in hypervisor state, and either address translation is enabled and \(\mathrm{VPM}_{1}=1\), or address translation is disabled, conditions that would have caused a Data Storage or an Instruction Storage interrupt if the affected memory were not virtualized instead cause a Hypervisor Data Storage or a Hypervisor Instruction Storage interrupt respectively. Because the Hypervisor Data Storage and Hypervisor Instruction Storage interrupts always put the thread in hypervisor state, they permit the hypervisor to handle the condition if appropriate (e.g., to restore the contents of a page that was reassigned), and to reflect it to the operating system's Data Storage or Instruction Storage interrupt handler otherwise.

When address translation is enabled, VPM mode has no effect on address translation. When address translation is disabled, addressing is controlled as specified in Section 5.7.3.

\subsection*{5.7.3 Hypervisor Real And Virtual Real Addressing Modes}

When a storage access is an instruction fetch performed when instruction address translation is disabled, or if the access is a data access and data address translation is disabled, it is said to be performed in "hypervisor real addressing mode" if the thread is in hypervisor state. If the thread is not in hypervisor state, the access is said to be performed in "virtual real addressing mode." Storage accesses in hypervisor real and virtual real addressing modes are performed in a manner that depends on the contents of I \(\mathrm{MSR}_{\mathrm{HV}}, \mathrm{VPM}\), PATE \(_{\text {PS }}\), HRMOR (see Chapter 2), bit 0 of the effective address (EAO), and the state of the Real Mode Storage Control Facility as described below. Bits \(1: 3\) of the effective address are ignored.
\(M_{S R}{ }_{H V}=1\)
- If \(E A_{0}=0\), the Hypervisor Offset Real Mode Address mechanism, described in Section 5.7.3.1, controls the access.
- If \(E A_{0}=1\), bits \(4: 63\) of the effective address are used as the real address for the access.

MSR \(_{\mathbf{H V}}=\mathbf{0}\)
- If PATE \(_{\text {HRIIGR }}=0 b 00\), the Virtual Real Mode Addressing mechanism, described in Section 5.7.3.3, controls the access.
- If PATE \(_{\text {HRIIGR }} \neq 0 \mathrm{bOO}\), partition-scoped translation is performed on the effective address. (See Section 5.7.12.3, "Obtaining Host Real Address, Radix on Radix".)

\subsection*{5.7.3.1 Hypervisor Offset Real Mode Address}

If \(\mathrm{MSR}_{\mathrm{HV}}=1\) and \(E A_{0}=0\), the access is controlled by the contents of the Hypervisor Real Mode Offset Register, as follows.

\section*{Hypervisor Real Mode Offset Register (HRMOR)}

Bits 4:63 of the effective address for the access are ORed with the 60-bit offset represented by the contents of the HRMOR, and the 60-bit result is used as the real address for the access. The supported offset values are all values of the form \(\mathrm{i} \times 2^{r}\), where \(0 \leq i<2^{j}\), and \(j\) and \(r\) are implementa-tion-dependent values having the properties that \(12 \leq r \leq 26\) (i.e., the minimum offset granularity is 4 KB and the maximum offset granularity is 64 MB ) and \(j+r=m\), where the real address size supported by the implementation is m bits.

\section*{Programming Note}
\(\mathrm{EA}_{4: 63-\mathrm{r}}\) should equal \({ }^{60-\mathrm{r}} 0\). If this condition is satisfied, ORing the effective address with the offset produces a result that is equivalent to adding the effective address and the offset.

If \(\mathrm{m}<60, \mathrm{EA}_{4: 63-\mathrm{m}}\) and \(\mathrm{HRMOR}_{4: 63-\mathrm{m}}\) must be zeros.

\subsection*{5.7.3.2 Storage Control Attributes for Accesses in Hypervisor Real Addressing Mode}

Storage accesses in hypervisor real addressing mode are performed as though all of storage had the following storage control attributes, except as modified by the Hypervisor Real Mode Storage Control facility (see Section 5.7.3.2.1). (The storage control attributes are defined in Book II.)
- not Write Through Required
- not Caching Inhibited, for instruction fetches
- not Caching Inhibited, for data accesses except those caused by the Load/Store Caching Inhibited instructions; Caching Inhibited, for data accesses caused by the Load/Store Caching Inhibited instructions
- Memory Coherence Required, for data accesses
- Guarded
- not SAO

Additionally, storage accesses in hypervisor real addressing mode are performed as though all storage was not No-execute.
\[
\begin{array}{l|l}
\text { I } & \begin{array}{l}
\text { Because storage accesses in hypervisor real } \\
\text { addressing mode do not use the SLB or the Page } \\
\text { I } \\
\text { Table, accesses in this mode bypass all checking } \\
\text { and recording of information contained therein } \\
\text { (e.g., storage protection checks that use informa- } \\
\text { tion contained therein are not performed, and refer- } \\
\text { ence and change information is not recorded). }
\end{array} .
\end{array}
\]

\subsection*{5.7.3.2.1 Hypervisor Real Mode Storage Control}

The Hypervisor Real Mode Storage Control facility provides a means of specifying portions of real storage that are treated as non-Guarded in hypervisor real addressing mode ( \(\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR}=0 \mathrm{~b} 10\), and \(\mathrm{MSR}_{\mathrm{IR}}=0\) or \(M_{S R}=0\), as appropriate for the type of access). The remaining portions are treated as Guarded in hypervisor real addressing mode. The means is a hypervisor resource (see Chapter 2), and may also be sys-tem-specific.
| The facility divides real storage into history blocks, in implementation-specific sizes. The history for instruction fetches is tracked separately from that for data accesses. If there is no instruction fetch history for a block and it is the target of an instruction fetch, the access is performed as though the block is Guarded, but the block is treated as non-Guarded for subsequent instruction fetches on a best effort basis, limited by the amount of history that the facility can maintain. If there is no data access history for a block and it is accessed using a Load/Store Caching Inhibited instruction, the access is performed as though the block is Guarded, and the block is treated as Guarded for subsequent accesses on a best effort basis, limited by the amount of history that the facility can maintain. If there is no data access history for a block and it is accessed using any other Load or Store instruction, the access is performed as though the block is Guarded, but the block is treated as non-Guarded for subsequent accesses on a best effort basis, limited by the amount of history that the facility can maintain.
The storage location specified by a Load/Store Caching Inhibited instruction must not be in storage that is specified by the Hypervisor Real Mode Storage Control
I facility to be treated as non-Guarded. The storage location specified by any other Load or Store instruction must not be in storage that is specified by the Hypervisor Real Mode Storage Control facility to be treated as
I Guarded. ("specified by the Hypervisor Real Mode Storage Control facility" means "specified in a history
| block".) The history can be erased using an slbia instruction; see Section 5.9.3.2.

\section*{Programming Note}

There are two cautions about mixing different types of accesses (i.e.Load/Store Caching Inhibited instructions vs. any other Load or Store instruction vs. instruction fetches). The first, as indicated above, is to avoid confusing the history mechanism, and the granularity for concern is a history block. For this caution, instruction fetches are irrelevant because they have their own history mechanism and are always intended to be non-guarded.

The second caution is to avoid storage paradoxes that result from a Caching Inhibited access to a location that is held in a cache. The nature of this caution and its solution are described in Section 5.8.2.2, "Altering the Storage Control Bits". The minimum granularity for concern is the history block, but may be larger, depending on extant translations to the storage in question. Since the consistency of instruction storage is managed by software and hypervisor real mode instruction fetches are always not Caching Inhibited, instruction fetches are also irrelevant to this caution.

The facility does not apply to implicit accesses to the Page Table performed during address translation or in recording reference and change information. These accesses are performed as described in Section 5.7.3.4.

\section*{Programming Note}

The preceding capability can be used to improve the performance of hypervisor software that runs in hypervisor real addressing mode, by causing accesses to instructions and data that occupy well-behaved storage to be treated as non-Guarded.

\subsection*{5.7.3.3 Virtual Real Mode Addressing Mechanism}

If \(\mathrm{MSR}_{\mathrm{HV}}=0\), the partition is using Paravirtualized HPT translation (PATE \(\mathrm{ERR}^{\prime} \| \mathrm{GR}=0 \mathrm{~b} 00\) ), and \(\mathrm{MSR}_{\mathrm{DR}}=0\) or \(M_{2 R}=0\) as appropriate for the type of access, the access is said to be made in virtual real addressing mode and is controlled by the mechanism specified below. The set of storage locations accessible by code is referred to as the Virtualized Real Mode Area (VRMA).
In virtual real addressing mode, address translation, storage protection, and reference and change recording are handled as follows.
- Address translation and storage protection are handled as if address translation were enabled, except that translation of effective addresses to virtual addresses use the SLBE values in Figure 18 instead of the entry in the SLB corresponding to the ESID. In this translation, bits 0:23 of the effec-
tive address are ignored (i.e., treated as if they were 0 s), bits 24:63-m may be ignored if \(m<40\), and the Virtual Page Class Key Protection mechanism does not apply.

\section*{- Programming Note}

The Virtual Page Class Key Protection mechanism does not apply because the authority mask that an OS has set for application programs executing with address translation enabled may not be the same as the authority mask required by the OS when address translation is disabled, such as when first entering an interrupt handler.
- Reference and change recording are handled as if address translation were enabled.
\begin{tabular}{|l|l|}
\hline Field & Value \\
\hline ESID & \({ }^{36} 0\) \\
\hline V & 1 \\
\hline B & Ob01-1 TB \\
\hline VSID & Ob00 II 0x0_01FF_FFFF \\
\hline \(\mathrm{K}_{\mathrm{s}}\) & 0 \\
\hline \(\mathrm{~K}_{\mathrm{p}}\) & undefined \\
\hline N & 0 \\
\hline L & PATE \(_{\text {PS[0] }}\) I \\
\hline C & 0 \\
\hline LP & PATE \(_{\text {PS[1:2] }}\) \\
\hline
\end{tabular}

Figure 18. SLBE for VRMA

\section*{Programming Note}

The \(C\) bit in Figure 18 is set to 0 because the imple-mentation-dependent lookaside information associated with the VRMA is expected to be long-lived. See Section 5.9.3.2.

\section*{Programming Note}

The 1 TB VSID 0x0_01FF_FFFF should not be used by the operating system for purposes other than mapping the VRMA when address translation is enabled.

\section*{Programming Note}

Software should specify \(\mathrm{PTE}_{\mathrm{B}}=0 \mathrm{~b} 01\) for all Page Table Entries that map the VRMA in order to be consistent with the values in Figure 18.

\section*{Programming Note}

All accesses to the RMA are considered not Guarded. The \(G\) bit of the associated Page Table Entry determines whether an access to the VRMA is Guarded. Therefore, if an instruction is fetched from the VRMA, a Hypervisor Instruction Storage interrupt will result if \(\mathrm{G}=1\) in the associated Page Table Entry.

\section*{Programming Note}

The RMA is considered non-SAO storage. However, any page in the VRMA is treated as SAO storage if WIMG \(=0 b 1110\) in the associated Page Table Entry.

\subsection*{5.7.3.4 Storage Control Attributes for Implicit Storage Accesses}
| Implicit accesses to a partition-scoped Page Table during address translation and in recording reference and change information are performed as though the storage occupied by the Page Table had the following storage control attributes.
- not Write Through Required
- not Caching Inhibited
- Memory Coherence Required
- not Guarded
- not SAO

Implicit accesses to a process-scoped Page Table during address translation and in recording reference and change information are performed using the storage control attributes in the partition-scoped Page Table that maps the process-scoped Page Table Entry that is being accessed.

\subsection*{5.7.4 Definitions}
process-scoped: Refers to translation performed using tables pointed to by Process Table Entries: guest Radix Tree translation in Radix on Radix mode, host Radix Tree translation in Radix on Radix when \(\mathrm{MSR}_{\mathrm{HV}}=1\) for quadrants 0 and 3, or Segment translation.
partition-scoped: Refers to translation performed using table(s) found using the first doubleword of Partition Table Entries, either host Radix Tree translation or HPT translation..
fully-qualified address: Refers to the address to be translated, when qualified by the effective LPID and effective PID.
guest real address: Refers to the input to the parti-tion-scoped translation process in Radix on Radix mode.
virtual address: Refers to the output of Segment translation and input to HPT translation.
host real address: Refers to the output of the parti-tion-scoped translation process when two levels of translation are being performed or the output of the pro-cess-scoped translation in Radix on Radix when \(\mathrm{MSR}_{\mathrm{HV}}=1\) for quadrants 0 and 3 . The simpler "real address" may be used interchangeably.

Page Directory: A table within the Radix Tree translation structure that contains elements ("Page Directory Entries") that point to other tables, instead of containing just Page Table Entries. The Page Directory that is at the root of the Radix Tree is called the "Root Page Directory."
effLPID, effPID: This is shorthand for effective LPID and effective PID. In certain circumstances, the value used for the LPID and/or the PID is specified to be zero instead of the actual register contents. "Effective" or "eff" is used to indicate the possibility of such a substitution. This value substitution happens only in Radix Tree translation, and is based on the value of \(E A_{0: 1}\) (see Section 5.7.5.1, "Effective Address Space Structure for Radix-using Partitions"). Value substitution does not happen in HPT translation. When a guest has its own Radix Tree (PATE \(\mathrm{GR}^{=1}\) ), PID substitution may take place. When a host uses Radix Tree translation (PATE HR \(=1\) ), both PID and LPID substitution may take place. When a host uses HPT translation, the only special significance associated with LPIDR \(=0\) is with regard to Segment Table walk when \(\mathrm{MSR}_{\mathrm{HV}}=1\), as described later.
adjunct: An adjunct is a software entity that resides in a partition along with an operating system and its applications in order to efficiently provide services (e.g. device drivers) for the partition. The adjunct partition is managed by the hypervisor. It runs in problem state with \(M S R_{H V ~ P R}=0 b 11\), thereby restricting the resources it can modify ( \(\mathrm{MSR}_{\mathrm{PR}}=1\) ) and causing its interrupts to go to the hypervisor ( \(\mathrm{MSR}_{\mathrm{HV}}=1\) ). It shares an HPT with the partition it serves. The adjunct is kept separate from the partition using Virtual Page Class Key protection. (The adjunct partition's lightness of weight derives from not requiring a full partition context switch (SLB flush, TLB flush, PTCR change, etc.) when the client partition invokes the services of the adjunct partition.) Each thread may have its own unique translations for an adjunct. As a result, adjunct segment descriptors cannot exist in the process's Segment Table and must instead be bolted in the SLB manually. The adjunct construct exists only with a hypervisor that uses HPT translation and only for LPID \(\neq 0\). The adjunct has its own 64-bit EA space. Accesses performed by the adjunct are not subject to guest translation (when using nested translation). Entry to an adjunct is only possible from hypervisor state. Prior to dispatching the adjunct, the hypervisor must invalidate SLB entries that map the effective address range that will be used by the adjunct.

Similarly, on exit from the adjunct, the hypervisor must invalidate its SLB entries

\subsection*{5.7.5 Address Ranges Having Defined Uses}

The address ranges described below have uses that are defined by the architecture.
- Fixed interrupt vectors

Except for the first 256 bytes, which are reserved for software use, the real page beginning at real address 0x0000_0000_0000_0000 is either used for interrupt vectors or reserved for future interrupt vectors.
- Implementation-specific use

The two contiguous real pages beginning at real address 0x0000_0000_0000_1000 are reserved for implementation-specific purposes.
- Offset Real Mode interrupt vectors

The real page beginning at the real address specified by the HRMOR is used similarly to the page for the fixed interrupt vectors.
- Relocated interrupt vectors

Depending on the values of \(M S R_{I R} D R\) and LPCR \(_{\text {AIL }}\) and on whether the specific interrupt will cause \(\mathrm{MSR}_{\mathrm{HV}}\) to change, either the virtual page containing the byte addressed by effective address 0x0000_0000_0001_8000 or the virtual page containing the byte addressed by effective address 0xC000_0000_0000_4000 may be used similarly to the page for the fixed interrupt vectors. (See Section 2.2.)
- System Call Vectored interrupt vectors

Depending on the value of LPCR \({ }_{\text {AIL }}\), the virtual page containing the effective address 0x0000_0000_0001_7000 or 0xc000_0000_0000_3000 contains the interrupt vectors that are invoked by the System Call Vectored instruction.
- Page Table

A contiguous sequence of real pages beginning at the real address specified by the first doubleword of the Partition Table Entry when HR=1 contains the Page Table.
- Adjunct Virtual Address Space

In systems in which the hypervisor uses HPT translation, the following virtual address ranges are reserved for adjunct use:
FFFFFD1FFF0000000 to FFFFFDFFFFFFFFFFFF,

FFFFFE1FFF0000000 to FFFFFEFFFFFFFFFFFF, and
FFFFFFOFFF0000000 to FFFFFFEFFFFFFFFFFF

\subsection*{5.7.5.1 Effective Address Space Structure for Radix-using Partitions}

When Radix Tree translation is in use but translation is off \(\quad\left(\mathrm{MSR}_{\mathrm{PR}}=0\right), \quad M S R_{\mathrm{HV}}\) selects between parti-tion-scoped translation of the real mode guest real address, formed by treating \(\mathrm{EA}_{0: 1}\) as 0 b 00 , and hypervisor real mode (see Section 5.7.3). When Radix Tree translation is in use and translation is on, \(\mathrm{EA}_{0: 1}\) together with \(M S R_{H V}\) are used to select one of as many as four Radix Trees with which to perform pro-cess-scoped translation, as a technique to make system calls and interrupts more efficient by avoiding the need to immediately change the contents of the PIDR and LPIDR. (See Figure 19 for an illustration of the mappings. Note that the second and third columns and related description below apply only when the HR field of the Partition Table Entry (PATE \({ }_{H R}\), see Figure 21) for the partition is set to 1. Also note that PATE \(_{\text {HR }}\) comes from the Partition Table Entry (PATE) for the partition indicated by the LPIDR, rather than effLPID.) Since there's nothing to prevent a process from generating any address in the 64b EA space, the exceptional cases are defined as follows. When a quadrant of the EA space has no associated Radix Tree, access to it results in an Instruction Segment exception or Data Segment exception, as appropriate for the type of access. Similarly, reference to any portion of these quadrants or the real mode guest real address described above that is not mapped by a Radix Tree (versus mapped by an invalid entry) will cause an Instruction or Data Segment exception.

\section*{- Programming Note}

Note that the quadrant structure is only available to software running in 64b mode. 32b software will only be able to access storage mapped by its own Radix Tree.

\section*{- Programming Note}

Warning: The functionality described in this section, e.g. directing most hypervisor interrupts to the LPID=0 translation tables, places great importance on the correctness of the format of and mappings in Partition Table Entry 0 and the tables it anchors. An error in any of these structures could have severe consequences including system checkstops and hangs.

\section*{Programming Note}

The intent is that the PIDR and LPIDR contents indicate the process and partition on behalf of which execution is taking place. For example, when a guest process interrupts to the hypervisor, execution to service the interrupt will generally be on behalf of the guest partition. When execution changes to be purely managing hypervisor resources that are not directly tied to any partition, the hypervisor should set LPIDR to 0 .

For guest and host applications and the guest operating system, quadrant \(0\left(\mathrm{EA}_{0: 1}=0 \mathrm{~b} 00\right)\) addresses the Radix Tree for the application and quadrant 3 \(\left(E A_{0: 1}=0 b 11\right)\) addresses the direct supervisor of the application. For the guest and host applications, it will frequently be the case that page protection is used to prevent access to quadrant 3 , but partition-wide shared text and/or data may also be located there. Quadrants 1 and 2 have no associated Radix Tree.

\section*{Programming Note}

Outboard accelerators may commonly be limited to accessing quadrants 0 and 3 as a matter of platform architecture. In such platforms, references to quadrants 1 and 2 may be regarded as errors.

For the hypervisor, quadrants 0 and 3 are as described above. Quadrant 1 ( \(\mathrm{EA}_{0: 1}=0 \mathrm{~b} 01\) ) addresses the guest application and quadrant \(2\left(\mathrm{EA}_{0: 1}=0 \mathrm{~b} 10\right)\) addresses the guest operating system, one of which experienced a hypervisor interrupt or performed a system call to the hypervisor. It will rarely be the case that quadrants 0 and 1 will be in use concurrently. A new value will usually be put in PIDR between accesses to quadrants 0 and 1.
When \(\mathrm{MSR}_{\mathrm{HV}}=1\) and \(E A_{0: 1}=0 b 00\) or \(0 b 11\), only pro-cess-scoped translation is performed. When \(\mathrm{MSR}_{\mathrm{HV}}=0\) and \(M S R_{I R / D R}=0\), only partition-scoped translation is performed. Otherwise, nested process- and parti-tion-scoped translations are performed.
\begin{tabular}{|c|}
\hline Guest \\
\hline \[
\begin{aligned}
& E A_{0: 1}=0 b 11 \\
& \text { effPID=0 } \\
& \text { effLPID=LPIDR }
\end{aligned}
\] \\
\hline \\
\hline \\
\hline \[
\begin{aligned}
& \text { EA }_{0: 1}=0 \mathrm{bOO} \\
& \text { effPID=PIDR } \\
& \text { effLPID=LPIDF }
\end{aligned}
\] \\
\hline
\end{tabular}
\begin{tabular}{|l|}
\multicolumn{1}{c}{ Host App } \\
\hline \begin{tabular}{l} 
EA \\
\(0: 1\) \\
effPID=0b11 \\
effLPID \(=0\)
\end{tabular} \\
\hline \\
\hline
\end{tabular}


Figure 19. Effective address space structure when using Radix Tree translation

\subsection*{5.7.6 In-Memory Tables}

The In-Memory Tables are used to find the tables that are used in the actual translation process for the partition and process that are executing. They enable hardware, including accelerator hardware separate and distinct from the Power ISA processors in the platform, to perform the translation process largely without software intervention. Description of the In-Memory Table structure follows. Hardware may cache the contents of the In-Memory Tables. Variants of tlbie[I] may be used to manage the caching even though the In-Memory Table contents are not cached in the TLB.

When an address in the In-Memory Table structure is specified to be a virtual or guest real address, the access to that address is considered to be performed with translation on. For a host using HPT translation, a base page size is specified for each such access to be used in the HPT search. The hypervisor can override the Segment Table Page Size in the Process Table Entry (PRTE base page size for the Process Table (PATE \({ }_{\text {PRTPS }}\) ) can be safely altered by the hypervisor since the OS does not have direct access to the Partition Table Entry. All accesses to the In-Memory Tables, the Segment Tables, and the guest Radix Tables that are performed with translation on, including for instruction address translation, are data accesses performed as if \(\mathrm{MSR}_{\mathrm{PR}}=0\) for the purpose of determing storage protection, although instruction side translation exceptions cause \([\mathrm{H}] \mathrm{ISI}\). (A specific example of the implications of
this is that tables used to translate instruction fetches may be located in guarded or no-execute storage.)

\section*{Programming Note}

The descriptors in the entries in this section contain addresses that are properly aligned so that no shifting is required. For example, the minimum size of the Partition Table is 4KB, so PATB has the thirteenth least significant address bit as its least significant bit. To construct the real address for a 4KB table, 12 zeros are appended on the right, and an appropriate number of address bits are removed from the left to match the real address size ( m ) supported by the implementation. For an 8 K table, bit 51 of the PTCR would be disregarded, and 13 zeros would be appended.

\subsection*{5.7.6.1 Partition Table}

The Partition Table Control Register (PTCR) is a hypervisor privileged SPR that contains the host real address of the base of the Partition Table and specifies its size. Software must ensure that the contents of the PTCR are the same for all processors in the system prior to enabling translation or transfering control to a partition..
\begin{tabular}{|c|cc|c|r|}
\hline\(/ / /\) & PATB & \(/ /\) & PATS \\
\hline 03 & 51 & 58 & 63 \\
\hline
\end{tabular}

\section*{Partition Descriptor}
\begin{tabular}{cll} 
Bit(s) & Name & Description \\
\(4: 51\) & PATB & Partition Table Base \\
\(59: 63\) & PATS & Partition Table Size \(=2^{12+\text { PATS }}\) \\
& & PATS \(\leq 24\)
\end{tabular}

All other fields are reserved.
Figure 20. Partition Table Control Register

> Programming Note
> If it becomes necessary to shrink the Partition Table or to change PATB to point to a table that is not identical to the existing one, it is necessary to issue tlbie with RIC=2 to invalidate caching of outdated In-Memory Table Entries.

The Parition Table is composed of a pair of doublewords per partition. The first doubleword indicates whether the host uses HPT or Radix Tree translation, and contains the base of the host's translation table structure in host real memory. The first doubleword also contains the size of the table structure and the size of the Root Page Directory for a hypervisor using Radix Tree translation, or the base page size for the VRMA for Paravirtualized HPT translation. Additional details about the parameters for HPT translation follow.

The HTABORG field contains the high-order 42 bits of the 60 -bit real address of the Page Table. The Page

Table is thus constrained to lie on a \(2^{18}\) byte ( 256 KB ) boundary at a minimum. At least 11 bits from the hash function (see Figure 29) are used to index into the Page Table. The minimum size Page Table is \(256 \mathrm{~KB}\left(2^{11}\right.\) PTEGs of 128 bytes each).
The Page Table can be any size \(2^{n}\) bytes where \(18 \leq n \leq 46\). As the table size is increased, more bits are used from the hash to index into the table.
\| The HTABSIZE field contains an integer giving the number of bits (in addition to the minimum of 11 bits) from the hash that are used in the Page Table index. This number must not exceed 28. HTABSIZE is used to generate a mask of the form 0b00...011...1, which is a string of 28 - HTABSIZE 0-bits followed by a string of HTABSIZE 1 -bits. The 1 -bits determine which additional bits (beyond the minimum of 11) from the hash are used in the index (see Figure 29).
On implementations that support a real address size of only m bits, \(\mathrm{m}<60\), bits \(0: 59-\mathrm{m}\) of the HTABORG field are treated as reserved bits, and software must set them to zeros.

\section*{Programming Note}

Let n equal the virtual address size (in bits) supported by the implementation. If \(n<67\), software should set the HTABSIZE field to a value that does not exceed n-39. Because the high-order 78-n bits of the VSID are assumed to be zeros, the hash value used in the Page Table search will have the high-order 67-n bits either all 0s (primary hash; see Section 5.7.9.2) or all 1 s (secondary hash). If HTABSIZE > n-39, some of these hash value bits will be used to index into the Page Table, with the result that certain PTEGs will not be searched.

\section*{Example:}

Suppose that the Page Table is \(16,384\left(2^{14}\right) 128\)-byte PTEGs, for a total size of \(2^{21}\) bytes ( 2 MB ). A 14-bit index is required. Eleven bits are provided from the hash to start with, so 3 additional bits from the hash must be selected. Thus the value in HTABSIZE must be 3 and the value in HTABORG must have its low-order 3 bits (bits \(43: 45\) of the first doubleword of the Partition Table Entry) equal to 0 . This means that the Page Table must begin on a \(2^{3+11+7}=2^{21}=2 \mathrm{MB}\) boundary.
\begin{tabular}{|c|c|c|c|c|c|}
0 & 3 & & 5 & 55 & 58 \\
63 \\
\hline 0 & \(/\) & HTABORG & \(/ /\) & & PS \\
\hline 0 & PRTB & & \(/ / /\) & & PRTPSSIZE \\
\hline 0 & \multicolumn{2}{c}{38} & & 55 & 58 \\
\hline
\end{tabular}

Paravirtualized HPT Partition Table Entry
\begin{tabular}{|c|c|c|}
\hline Bit(s) & \multirow[t]{6}{*}{Name
HR} & Description \\
\hline \multirow[t]{5}{*}{0} & & Host Radix \\
\hline & & Ob0- hypervisor uses HPT \\
\hline & & translation for this partition \\
\hline & & Ob1- hypervisor uses Radix \\
\hline & & Tree translation for this partition \\
\hline 4:45 & HTABORG & Hashed Page Table Base \\
\hline 56:58 & PS & Page Size (uses LIILP encoding as in current SLBE) \\
\hline \multirow[t]{2}{*}{59:63} & \multirow[t]{2}{*}{HTABSIZE} & HPT size \(=2^{\text {HTABSIZE+18 }}\) \\
\hline & & HTABSIZE \(\leq 28\) \\
\hline \multirow[t]{3}{*}{0} & \multirow[t]{3}{*}{GR} & Guest Radix \\
\hline & & Ob0- partition uses HPT \\
\hline & & Ob1-partition uses Radix Tree \\
\hline 1:38 & PRTB & Process Table Base (when UPRT=1) \\
\hline 56:58 & PRTPS & Process Table Page Slze (when UPRT=1) (uses LIILP encoding as in current SLBE) \\
\hline 59:63 & PRTS & Process Table Size \(=2^{12+\text { PRTS }}\) PRTS \(\leq 24\) (when UPRT=1) \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|c|c|c|}
\hline 0 & \multicolumn{2}{|r|}{3} & & 58 & 63 \\
\hline 1 & RTS1 & / & RPDB & RTS2 & RPDS \\
\hline 1 & / & & PRTB & // & PRTS \\
\hline 0 & & & & 58 & 63 \\
\hline
\end{tabular}

\section*{Radix on Radix Partition Table Entry}
\begin{tabular}{|c|c|c|}
\hline Bit(s) & Name & Description \\
\hline 0 & HR & Host Radix \\
\hline & & Ob0- hypervisor uses HPT translation for this partition \\
\hline & & Ob1- hypervisor uses Radix Tree translation for this partition \\
\hline 1:2 & RTS1 & Radix Tree Size[0:1] \\
\hline 4:55 & RPDB & Root Page Directory Base \\
\hline 56:58 & RTS2 & Radix Tree Size[2:4] (number of address bits mapped), size \(=2^{\text {RTS }+31}\) \\
\hline 59:63 & RPDS & Root Page Directory Size \(=2^{\text {RPDS }+3}\), RPDS \(\geq 5\) \\
\hline 0 & GR & Guest Radix \\
\hline & & \begin{tabular}{l}
Ob0- partition uses HPT \\
Ob1- partition uses Radix Tree
\end{tabular} \\
\hline 4:51 & PRTB & Process Table Base \\
\hline 59:63 & PRTS & Process Table Size \(=2^{12+\text { PRTS }}\) \\
\hline & & PRTS \(\leq 24\) (when UPRT=1) \\
\hline
\end{tabular}

All other fields are reserved.
Figure 21. Partition Table Entry Variants
The second doubleword of the Partition Table Entry indicates whether the guest has its own Radix Tree. It also contains the base of the partition's Process Table, which is a guest real address (or effective address when effective LPID=0) for radix hypervisor and virtual address for HPT hypervisor, and the size of the Pro-
cess Table. When Segment Tables are provided, the Process Table base address is specified as a VSID with the assumption that the Process Table is located at zero offset in the segment, and also includes the base page size used for the HPT search, with the rest of the implied segment descriptor being \(\mathrm{B}=0 \mathrm{~b} 01\) (1TB segment), \(K s=K p=0, N=0, C=0\), and virtual page class key protection does not apply. The Partition Table Entry variants with \(\mathrm{HR} \neq \mathrm{GR}\) are reserved, and will be reported as an unsupported MMU configuration type of HDSI or HISI. Other variants are illustrated in Figure 21.Notethat a configuration with \(\mathrm{HR}=1\) for a non-zero LPID and HR=0 for LPID=0 is considered an unsupported MMU configuration because it would attempt to perform HPT translation in quadrants 0 and 3 when \(\mathrm{MSR}_{\mathrm{HV}}=1\). In addition, LPID=0 with Radix Tree translation is an unsupported MMU configuration when \(\mathrm{MSR}_{\mathrm{HV}}=0\).

\begin{abstract}
- Programming Note

The sizes of the Partition and Process Tables are provided to simplify hardware design and testing. The size enables the hardware to mask address bits instead of providing an adder. No size checking is provided for these tables. (An out-of-range LPID or PID will not produce an exception simply because of its size.) Hypervisor software may protect against such errors by the OS by not providing a translation for virtual / guest real addresses beyond the end of the Process Table.
\end{abstract}

\subsection*{5.7.6.2 Process Table}

The Process Table is composed of a quadword Process Table Entry per process in the partition. For partitions that use HPT translation ( \(\mathrm{HR}=0\) and \(\mathrm{GR}=0\) ), the Process Table Entry contains a Segment Table descriptor, which is composed of the origin of the Segment Table in virtual address space, the size of the segment and pages that hold the table, the size of the table, and a valid bit that is turned off while changes are made to the entry and Segment Table. The translation of the base address of the Segment Table is completed using an implied segment descriptor with \(\mathrm{Ks}=\mathrm{Kp}=0, \mathrm{~N}=0\), \(\mathrm{C}=0\), and virtual page class key protection does not apply. For partitions that use Radix Tree translation, the Process Table Entry contains a Radix Tree root descriptor. When running on a host that uses Radix Tree translation, there are two cases. When effLPID=0, the RPDB is a host real address. Otherwise, the address is a guest real address and must undergo translation using the hypervisor's Radix Tree for the partition (i.e. the "partition-scoped" tables, as defined later).
i
| \(0 \quad 1\) 63
\begin{tabular}{|c|c|c|c|c|}
\hline B & \multicolumn{4}{|c|}{ STABORGU } \\
\hline STABORGL & I/I & STABSIZE & STPS & V \\
\hline 0 & 3 & & 55 & 59 \\
\hline
\end{tabular}

```

DW Bit(s) Name Description
0 1:2 RTS1 Radix Tree Size[0:1]
3 / Reserved
4:55 RPDB Root Page Directory Base
56:58 RTS2 Radix Tree Size[2:4] (number of
address bits mapped),
size=2}\mp@subsup{2}{}{\mathrm{ RTS+31}
59:63 RPDS Root Page Directory Size
=2}\mp@subsup{2}{}{\mathrm{ RPSS+3}},\textrm{RPDS}\geq

```

All other fields are reserved.
Figure 22. Process Table Entry Variants

\subsection*{5.7.7 Address Translation Overview}

The effective address (EA) is the address generated by the hardware for an instruction fetch or for a data access. If address translation is enabled, this address is passed to the Address Translation mechanism, which attempts to convert the address to a real address which
I is then used to access storage. If the effective address cannot be translated, a storage exception (see Section 5.2) occurs.

The architecture defines segment translation and two types of page translation. Segment translation is paired with HPT translation. The other supported "pairing" is two level Radix Tree translation. Either of these pairings can be used to translate an effective address into a host real address. The In-Memory Tables described above determine the translation mode used by a partition, as well as the locations of the Page Tables and Segment Tables, and the base page size for
the Segment Tables. When \(\mathrm{MSR}_{\mathrm{HV}}=1\) and/or \(\mathrm{MSR}_{\mathrm{IR}}=0\) or \(\mathrm{MSR}_{\mathrm{DR}}=0\) (as appropriate for the type of access), the steps taken for a given mode vary. See Sections 5.7.12.3 and 5.7.12.4 for details.

The pairing of Segment translation and Hashed Page Table (HPT) translation applies Segment translation to an effective address to produce a virtual address as described in Section 5.7.8, and HPT translation to the virtual address to produce a host real address as described in Section 5.7.9. The segment translation is managed cooperatively by the guest and the hypervisor, but the HPT translation is always managed by the hypervisor with the guest typically giving direction via system calls to the hypervisor in a paravirtualization relationship. This mode is commonly referred to as Paravirtualized HPT translation. The segment translation is managed on a per-process ("process-scoped") basis, mapping a smaller effective address space into a large "partition-scoped" virtual address space, where the segment can be used as a shared memory object. There is also the possibility of thread-unique mappings. In the basic version of HPT translation, storage exceptions are directed to the operating system, which in turn issues system calls to the hypervisor. When Virtualized Partition Memory is enabled, storage exceptions are directed to the hypervisor, enabling a higher degree of memory overcommitment as the hypervisor transparently steals pages from the partition. Figure 23 gives an overview of the address translation process.


Figure 23. Address translation overview
In Paravirtualized HPT mode, the hypervisor also uses the segment/HPT pairing, and can create a process called an "adjunct". To do so, it eliminates any potentially conflicting guest segment mappings and creates adjunct mappings prior to dispatching the adjunct.

In the other pairing, Radix Tree translation is used for both the process-scoped and partition-scoped mappings. This mode is commonly referred to as Radix on Radix translation. Figure 24 gives an overview of the address translation process for Radix on Radix translation. Note that each level of the guest Radix Tree produces a guest real address that must itself undergo partition-scoped translation. See Figure 35 for a detailed illustration of the entire process.

In Radix on Radix translation, storage exceptions for the process-scoped mappings are directed to the operating system, and storage exceptions for parti-tion-scoped mappings are directed to the hypervisor. As a result, the hypervisor can use the partition-scoped mapping to limit the size of the guest real address space, and Virtualized Partition Memory is not necessary to enable a higher degree of memory overcommitment. If in Radix on Radix mode the guest real address is outside the range covered by the partition-scoped Radix Tree, the results are boundedly undefined.

\section*{Programming Note}

Choices for GRA outside the parititon-scoped Radix Tree include ignoring the address bits that exceed the reach of the tree and causing an interrupt. Software must not depend on the bits being ignored because that could prevent future expansion of the GRA space.


Figure 24. Address translation overview, Radix on Radix

\section*{Translation Lookaside Buffer}

Conceptually, the Page Table is searched by the address relocation hardware to translate every reference. For performance reasons, the hardware usually keeps a Translation Lookaside Buffer (TLB) that holds PTEs that have recently been used. The TLB is searched prior to searching the Page Table. As a consequence, when software makes changes to the Page Table it must perform the appropriate TLB invalidate operations to maintain the consistency of the TLB with the Page Table (see Section 5.10).

An implementation may associate each of its TLB entries with the partition for which the TLB entry was created, so that the entries can be retained while other
partitions are executing. In this case, when a valid TLB entry is created, the LPID value from LPIDR is written into the TLB entry.

\section*{Programming Notes}
1. Page Table Entries may or may not be cached in a TLB.
2. It is possible that the hardware implements more than one TLB, such as one for data and one for instructions. In this case the size and shape of the TLBs may differ, as may the values contained therein.
3. Use the tlbie instruction to ensure that the TLB no longer contains a mapping for a particular virtual page.

\subsection*{5.7.8 Segment Translation}

For explicit accesses in Paravirtualized HPT mode, conversion of a 64-bit effective address to a virtual address is done by searching the Segment Lookaside Buffer (SLB) as shown in Figure 25. If no matching translation is found in the SLB, LPCR \({ }_{U P R T}=1\), and \(\mathrm{MSR}_{\mathrm{HV}}=0\) or in Paravirtualized HPT mode with PID=0, the Segment Table is searched. For implicit accesses, implicit segment descriptors are provided, as described elsewhere in this chapter.

64-bit Effective Address

Figure 25. Translation of 64-bit effective address to 78 bit virtual address

\subsection*{5.7.8.1 Segment Lookaside Buffer (SLB)}

The Segment Lookaside Buffer (SLB) specifies the mapping between Effective Segment IDs (ESIDs) and Virtual Segment IDs (VSIDs). The number of SLB entries is implementation-dependent, except that all implementations provide at least 32 entries.

The first four entries, and when LPCR \({ }_{\text {UPRT }}=0\) all of the entries, of the SLB are managed by software, using the instructions described in Section 5.9.3.2. See Chapter 11. "Synchronization Requirements for Context Alterations" on page 1127 for the rules that software must follow when updating the SLB.


\section*{SLB Entry}

Each SLB entry (SLBE, sometimes referred to as a "segment descriptor") maps one ESID to one VSID. Figure 26 shows the layout of an SLB entry
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline ESID & V & B & & VSID & \(\mathrm{K}_{\mathrm{s}} \mathrm{K}_{\mathrm{p}}\) NLC & / \\
\hline
\end{tabular}
\begin{tabular}{cll} 
Bit(s) & Name & Description \\
\(0: 35\) & ESID & Effective Segment ID \\
36 & V & Entry valid (V=1) or invalid (V=0) \\
\(37: 38\) & B & Segment Size Selector \\
& & Ob00 - 256 MB (s=28) \\
& & Ob01 - 1 TB (s=40) \\
& & Ob10 - reserved \\
Ob11 - reserved \\
\(39: 88\) & VSID & Virtual Segment ID \\
89 & \(\mathrm{~K}_{\mathrm{s}}\) & Supervisor (privileged) state stor- \\
& & age key (see Section 5.7.14.2) \\
90 & \(\mathrm{~K}_{\mathrm{p}}\) & Problem state storage key (See \\
& & Section 5.7.14.2.) \\
91 & N & No-execute segment if N=1 \\
92 & L & Virtual page size selector bit 0 \\
93 & C & Class \\
\(95: 96\) & LP & Virtual page size selector bits 1:2
\end{tabular}

All other fields are reserved. \(\mathrm{B}_{0}\left(\mathrm{SLBE}_{37}\right)\) is treated as a reserved field.

Figure 26. SLB Entry
Instructions cannot be executed from a No-execute ( \(\mathrm{N}=1\) ) segment.
Segments may contain a mixture of page sizes. The L and LP bits specify the base virtual page size for the segment. The SLB \({ }_{\text {LIILP }}\) encodings are those shown in Figure 27. The base virtual page size (also referred to as the "base page size") is the smallest virtual page size that can be used to map a given access, and in most cases is the smallest virtual page size for the segment. (The exception is that multiple base virtual page sizes can occur within the same segment when the base page size specified for a given implicit access (e.g. of one segment table) does not match the base page size specified for another implicit access (e.g. of a different segment table or the process table) or for explicit accesses. References to the base page size for a segment will be understood not to preclude or functionally conflict with this possibility.) The base virtual page size is \(2^{b}\) bytes. The actual virtual page size (also referred to as the "actual page size" or "virtual page size") is specified by PTE L LP
\begin{tabular}{|c|c|}
\hline encoding & base page size \\
\hline Ob000 & 4 KB \\
\hline Ob101 & 64 KB \\
\hline \begin{tabular}{c} 
additional \\
values
\end{tabular} \\
\hline The "additional values" are implementation-depen- \\
\begin{tabular}{c}
\(2^{b}\) bytes, where \(\mathrm{b}>12\) and b may differ \\
dent, as are the corresponding base virtual page \\
sizes. Any values that are not supported by a given \\
implementation are reserved in that implementa- \\
tion.
\end{tabular} \\
\hline
\end{tabular}

\section*{Figure 27. Page Size Encodings}

For each SLB entry, software must ensure the following requirements are satisfied.
- LIILP contains a value supported by the implementation.
- The base virtual page size selected by the L and LP fields does not exceed the segment size selected by the B field.
- If \(s=40\), the following bits of the SLB entry contain 0 s.
\[
\begin{array}{ll}
- & \text { ESID }_{24: 35} \\
\text { - } & \text { VSID }_{38: 49}
\end{array}
\]

The bits in the above two items are ignored by the hardware.
The Class field of the SLBE is used in conjunction with the slbie and slbia instructions (see Section 5.9.3.2). "Class" refers to a grouping of SLB entries and imple-mentation-specific lookaside information so that only entries in a certain group need be invalidated and others might be preserved. The Class value assigned to an implementation-specific lookaside entry derived from an SLB entry must match the Class value of that SLB entry. The Class value assigned to an implementa-tion-specific lookaside entry that is not derived from an SLB entry (such as real mode address "translations") is 0.

Software must ensure that the SLB contains at most one entry that translates a given effective address, and that if the SLB contains an entry that translates a given effective address, then any previously existing translation of that effective address has been invalidated. An attempt to create an SLB entry that violates this requirement may cause a Machine Check.

\section*{Programming Note}

It is permissible for software to replace the contents of a valid SLB entry without invalidating the translation specified by that entry provided the specified restrictions are followed. See Chapter 11 Note 10.

\subsection*{5.7.8.2 SLB Search}

When the hardware searches the SLB, all entries are tested for a match with the EA. For a match to exist, the following conditions must be satisfied for indicated fields in the SLBE.
- \(\mathrm{V}=1\)
- \(E S I D_{0: 63-\mathrm{s}}=\mathrm{EA}_{0: 63-\mathrm{s}}\), where the value of s is specified by the \(B\) field in the SLBE being tested
If no match is found, the search fails. If one match is found, the search succeeds. If more than one match is found, one of the matching entries is used as if it were the only matching entry, or a Machine Check occurs.
If the SLB search succeeds, the virtual address (VA) is formed from the EA and the matching SLB entry fields as follows.
\[
\text { VA=VSID }{ }_{0: 77-\mathrm{s}} \text { II EA } 64-\mathrm{s}: 63
\]

The Virtual Page Number (VPN) is bits 0:77-p of the virtual address. The value of \(p\) is the actual virtual page size specified by the PTE used to translate the virtual address (see Section 5.7.9.1). If \(\mathrm{SLBE}_{\mathrm{N}}=1\), the N (No-execute) value used for the storage access is 1 .

If the SLB search fails and the state is not such that a Segment Table search will be performed, a segment fault occurs. This is an Instruction Segment exception or a Data Segment exception, depending on whether the effective address is for an instruction fetch or for a data access.

\subsection*{5.7.8.3 Segment Table Description and Search}

The Segment Table is an aligned structure composed of 16B segment descriptors organized into 128 byte Segment Table Entry Groups (STEGs). Let q = STABSIZE+12, \(\log _{2}\) (size of the Segment Table). The base of the Segment Table in virtual address space is STABORG \({ }_{0: 77-q} \|^{9} 0\). Primary and secondary hashes are defined for 256 MB and 1TB segments, each mapping the ESID to an STEG. The appropriate number (for the size of the Segment Table) of low order ESID bits (their inverse, for the secondary hash) directly select the STEG. The order of STEG specification in the following subsections is the preferred order for a serial search. Implementations may search the STEGs in parallel. If no match is found, a segment fault occurs. If a serial search is done, the search may stop when a match has been found. If more than one match is found, one of the matching entries is used as if it were the only matching entry.
\begin{tabular}{|l|l|l|l|r|r|r|r|r|}
\hline ESID & V & // & B & VSID & \(\mathrm{K}_{\mathrm{s}} \mathrm{K}_{\mathrm{p}}\) NLC & / & LP & SW \\
\hline 0 & 35 & 36 & 63 & 65 & \multicolumn{2}{|c|}{115} & 120 & 121 \\
\hline
\end{tabular}
\begin{tabular}{|c|c|c|}
\hline Bit(s) & Name & Description \\
\hline 36 & V & Entry valid (V=1) or invalid (V=0) \\
\hline 64:65 & B & Segment Size Selector \\
\hline & & Ob00-256 MB ( \(\mathrm{s}=28\) ) \\
\hline & & 0b01-1 TB ( \(\mathrm{s}=40\) ) \\
\hline & & Ob10-reserved \\
\hline & & Ob11-reserved \\
\hline 66:115 & VSID & Virtual Segment ID \\
\hline 116 & \(\mathrm{K}_{\mathrm{s}}\) & Supervisor (privileged) state storage key (see Section 5.7.14.2) \\
\hline 117 & \(\mathrm{K}_{\mathrm{p}}\) & Problem state storage key (See Section 5.7.14.2.) \\
\hline 118 & N & No-execute segment if \(\mathrm{N}=1\) \\
\hline 119 & L & Virtual page size selector bit 0 \\
\hline 120 & C & Class \\
\hline 122:123 & LP & Virtual page size selector bits 1:2 \\
\hline 124:127 & SW & available for software use \\
\hline
\end{tabular}

All other fields are reserved.
Figure 28. Segment Table Entry

\subsection*{5.7.8.3.1 Primary Hash for 256MB Segment}

The STEG is located at host VA STABORG \(_{0: 77-q} \| E_{43-q: 35}\) II 0b0000000.
Each of the 8 SSTEs are searched to find a valid entry \((\mathrm{V}=1, \mathrm{~B}=0 \mathrm{~b} 00)\) that matches the ESID \(\left(\mathrm{STE}_{\mathrm{ESID}[0: 35]}=\right.\) \(E A_{0: 35}\) ) of the access being translated.

\subsection*{5.7.8.3.2 Primary Hash for 1TB Segment}

The STEG is located at host VA STABORG \(_{0: 77-q}\left\|E_{31-q: 23}\right\| 0^{0 b 0000000 .}\)
Each of the 8 SSTEs are searched to find a valid entry \((\mathrm{V}=1, \mathrm{~B}=0 \mathrm{Ob} 01)\) that matches the ESID (STE \(\mathrm{ESIID}[0: 23]=\) \(\left.E A_{0: 23}\right)\) of the access being translated.

\subsection*{5.7.8.3.3 Secondary Hash for 256MB Segment}

The STEG is located at host VA
STABORG \(_{0: 77-q}\left\|\neg \mathrm{EA}_{43-q: 35}\right\| \mathrm{Ob} 0000000\).
Each of the 8 SSTEs are searched to find a valid entry \((\mathrm{V}=1, \mathrm{~B}=0 \mathrm{~b} 00)\) that matches the ESID (STE \(\mathrm{ESID}[0: 35]=\) \(\left.E A_{0: 35}\right)\) of the access being translated.

\subsection*{5.7.8.3.4 Secondary Hash for 1TB Segment}

The STEG is located at host VA
STABORG \(_{0: 77-\mathrm{q}} \| \neg \mathrm{EA}_{31-q: 23}\) II 0b0000000.
Each of the 8 SSTEs are searched to find a valid entry \((\mathrm{V}=1, \mathrm{~B}=0 \mathrm{0b01})\) that matches the ESID (STE \(\mathrm{ESIID}[0: 23]=\) \(\left.E A_{0: 23}\right)\) of the access being translated.

\subsection*{5.7.9 Hashed Page Table Translation}

In Paravirtualized HPT mode, conversion of a 78-bit virtual address to a real address is done by searching the Page Table as shown in Figure 29.


Figure 29. Translation of 78-bit virtual address to 60-bit real address

\subsection*{5.7.9.1 Hashed Page Table}

The Hashed Page Table (HTAB) is a variable-sized data structure that specifies the mapping between virtual page numbers and real page numbers, where the real page number of a real page is bits 0:47 of the address of the first byte in the real page. The HTAB's size can be any size \(2^{n}\) bytes where \(18 \leq n \leq 46\). The HTAB must be located in storage having the storage control attributes that are used for implicit accesses to it (see Section 5.7.3.4). The starting address must be a multiple of \(2^{18}\) bytes.
The HTAB contains Page Table Entry Groups (PTEGs). A PTEG contains 8 Page Table Entries (PTEs) of 16 bytes each; each PTEG is thus 128 bytes long. PTEGs are entry points for searches of the Page Table.

See Section 5.10 for the rules that software must follow when updating the Page Table.

\section*{Programming Note}

The Page Table must be treated as a hypervisor resource (see Chapter 2), and therefore must be placed in real storage to which only the hypervisor has write access. Moreover, the contents of the Page Table must be such that non-hypervisor software cannot modify storage that contains hypervisor programs or data.

\section*{Page Table Entry}

Each Page Table Entry (PTE) maps one VPN to one RPN. Figure 30 shows the layout of a PTE. This layout is independent of the Endian mode of the thread.

\begin{tabular}{|c|c|c|c|}
\hline \multirow[t]{6}{*}{Dword 0} & Bit(s) & Name & Description \\
\hline & 12:56 & AVA & Abbreviated Virtual Address \\
\hline & 57:60 & SW & Available for software use \\
\hline & 61 & L & Virtual page size ObO-4 KB Ob1 - greater than 4KB (large page) \\
\hline & 62 & H & Hash function identifier \\
\hline & 63 & V & Entry valid (V=1) or invalid
\[
(\mathrm{V}=0)
\] \\
\hline \multirow[t]{3}{*}{1} & 0 & pp & Page Protection bit 0 \\
\hline & 2:3 & key & KEY bits 0:1 \\
\hline & 4:5 & B & \begin{tabular}{l}
Segment Size
Ob00-256 MB \\
Ob01-1 TB \\
Ob10 - reserved \\
Ob11 - reserved
\end{tabular} \\
\hline
\end{tabular}
\begin{tabular}{llll} 
Dword & \begin{tabular}{lll} 
Bit(s) & Name \\
\(7: 43\) & ARPN
\end{tabular} & \begin{tabular}{l} 
Description \\
Abbreviated Real Page \\
Number
\end{tabular} \\
& & & Large page size selector \\
\(44: 51\) & LP & Ler \\
\(52: 54\) & key & KEY bits \(2: 4\) \\
55 & R & Reference bit \\
56 & C & Change bit \\
\(57: 60\) & WIMG & Storage control bits \\
61 & N & No-execute page if N=1 \\
\(62: 63\) & pp & Page Protection bits \(1: 2\)
\end{tabular}

All other fields are reserved.
Figure 30. Page Table Entry

\section*{Programming Note}

The H bit in the Page Table Entry should not be set to one unless the secondary Page Table search has been enabled.

If \(b \leq 23\), the Abbreviated Virtual Address (AVA) field contains bits 0:54 of the VA. Otherwise bits 0:77-b of the AVA field contain bits \(0: 77-\mathrm{b}\) of the VA, and bits 78-b:54 of the AVA field must be zero.

\section*{Programming Note}

The AVA field omits the low-order 23 bits of the VA. These bits are not needed in the PTE, because the low-order b of these bits are part of the byte offset into the virtual page and, if \(b<23\), the high-order 23-b of these bits are always used in selecting the PTEGs to be searched (see Section 5.7.9.2).

On implementations that support a virtual address size of only \(n\) bits, \(n<78\), bits \(0: 77-n\) of the AVA field must be zeros.
A virtual page is mapped to a sequence of \(2^{p-12}\) contiguous real pages such that the low-order p-12 bits of the real page number of the first real page in the sequence are 0 s .

PTE \({ }_{\text {L }}\) specify both a base virtual page size (henceforth referred to as the "base page size") and an actual virtual page size (henceforth referred to as the "actual page size" or "virtual page size"). The actual page size is the size of the virtual page mapped by the PTE. The base page size is the smallest actual page size that a segment can contain for explicit accesses or for a given implicit access, and plays a role in the placement of the PTE in the HPT.
If \(P T E_{L}=0\), the base virtual page size and actual virtual page size are 4 KB , and ARPN concatenated with LP (ARPNIILP) contains the page number of the real page that maps the virtual page described by the entry.
If \(\mathrm{PTE}_{\mathrm{L}}=1\), the base page size and actual page size are specified by PTE \(_{\text {LP }}\) In this case, the contents of PTE \(_{\text {LP }}\) have the format shown in Figure 31. Bits labelled " \(r\) " are
bits of the real page number. Bits labelled "z" specify the base page size and actual page size. The values of the " \(z\) " bits used to specify each size are implemen-tation-dependent. The values of the " \(z\) " bits used to specify each size, along with all possible values of "r" bits in the LP field, must result in LP values distinct from other LP values for other sizes. Actual page sizes 4 KB and 64 KB are always supported; other actual page sizes are implementation-dependent. If \(\mathrm{PTE}_{\mathrm{L}}=1\), the actual page size must be greater than 4 KB . Which combinations of different base page size and actual page size are supported is implementation-dependent, except that the combination of a base page size of 4 KB with an actual page size of 64 KB is always supported.
\begin{tabular}{|c|c|}
\hline PTE LP & ac \\
\hline rrrr_rrrz & \(\geq 8 \mathrm{~KB}\) \\
\hline rrrr_rrzz & \(\geq 16 \mathrm{~KB}\) \\
\hline rrrr_rzzz & \(\geq 32 \mathrm{~KB}\) \\
\hline rrrr_zzzz & \(\geq 64 \mathrm{~KB}\) \\
\hline rrrz_zzzz & \(\geq 128\) KB \\
\hline rrzz_zzzz & \(\geq 256\) KB \\
\hline rzzz_zzzz & \(\geq 512 \mathrm{~KB}\) \\
\hline zzzz_zzzz & \(\geq 1 \mathrm{MB}\) \\
\hline
\end{tabular}

Figure 31. Format of \(\mathrm{PTE}_{\mathrm{LP}}\) when \(\mathrm{PTE}_{\mathrm{L}}=1\)
There are at least 2 formats of PTE 64 KB page. One format is used with \(\mathrm{SLBE}_{\text {LIILP }}=\) Ob000 and one format is used with \(\mathrm{SLBE}_{\text {LIILP }}=0 \mathrm{~b} 101\).

The actual page size selected by the LP field must not exceed the segment size selected by the B field. Forms of PTE \({ }_{\mathrm{LP}}\) not supported by a given implementation are treated as reserved values for that implementation.

The concatenation of the ARPN field and bits labeled " \(r\) " in the LP field contain the high-order bits of the real page number of the real page that maps the first 4 KB of the virtual page described by the entry.

The low-order \(\mathrm{p}-12\) bits of the real page number contained in the ARPN and LP fields must be Os and are ignored by the hardware.

\section*{Programming Note}

The actual page size specified by a given \(P T E_{L P}\) format is at least \(2^{12+(8-c)}\), where \(c\) is the number of \(r\) bits in the format.

\section*{Programming Note}

Implementations often have TLBs and implementa-tion-dependent lookaside buffers (e.g. ERATs) used to cache translations of recently used storage addresses. Mapping virtual storage to large pages may increase the effectiveness of such lookaside buffers, improving performance, because it is possible for such buffers to translate a larger range of addresses, reducing the frequency that the Page Table must be searched to translate an address.

Instructions cannot be executed from a No-execute ( \(\mathrm{N}=1\) ) page.

\section*{Page Table Size}

The number of entries in the Page Table directly affects performance because it influences the hit ratio in the Page Table and thus the rate of page faults. If the table is too small, it is possible that not all the virtual pages that actually have real pages assigned can be mapped via the Page Table. This can happen if too many hash collisions occur and there are more than 16 entries for the same primary/secondary pair of PTEGs (when the secondary Page Table search is enabled) or more than 8 entries for the same primary PTEG (when the secondary Page Table search is disabled).
While this situation cannot be guaranteed not to occur for any size Page Table, making the Page Table larger than the minimum size (see Section 5.7.6.1) will reduce the frequency of occurrence of such collisions.

\section*{Programming Note}

If large pages are not used, it is recommended that the number of PTEGs in the Page Table be at least half the number of real pages to be accessed. For example, if the amount of real storage to be accessed is \(2^{31}\) bytes ( 2 GB ), then we have \(2^{31-12}=2^{19}\) real pages. The minimum recommended Page Table size would be \(2^{18}\) PTEGs, or \(2^{25}\) bytes ( 32 MB ).

\subsection*{5.7.9.2 Page Table Search}

When the hardware searches the Page Table, the accesses are performed as described in Section 5.7.3.4.
An outline of the HTAB search process is shown in Figure 29. Up to two hash functions are used to locate a PTE that may translate the given virtual address.
1. A 39-bit hash value is computed from the VA. The value of \(s\) is the value specified in the SLBE that was used to generate the virtual address; the value of \(b\) is equal to \(\log _{2}\) (base page size specified in the SLBE that was used to translate the address).Primary Hash:
If \(s=28\), the hash value is computed by Exclusive ORing \(\mathrm{VA}_{11: 49}\) with ( \({ }^{11+\mathrm{b}} \mathrm{O}_{\mathrm{IIVA}}^{50: 77-\mathrm{b}}\) )
If \(s=40\), the hash value is computed by Exclusive ORing the following three quantities: \(\left(\mathrm{VA}_{24: 37}\right.\) \(\left.\|{ }^{25} 0\right)\), ( \(0 \| V A_{0: 37}\) ), and ( \(\left.{ }^{b-1} 0 \| V A_{38: 77-b}\right)\)
The 60-bit real address of a PTEG is formed by concatenating the following values:
■ Bits 0:27 of the 39-bit appropriate primary or secondary hash value ANDed with the mask generated from bits 59:63 of the first double-
word of the Partition Table Entry (HTABSIZE) and then added to the value of bits \(4: 45\) of the first doubleword of the Partition Table Entry (HTABORG).
- Bits 28:38 of the 39 -bit hash value.

■ Seven 0-bits.
This operation identifies a particular PTEG, called the "primary PTEG", whose eight PTEs will be tested.

\section*{2. Secondary Hash:}

If the secondary Page Table search is enabled ( \(\mathrm{LPCR}_{\mathrm{TC}}=0\) ), perform the secondary hash function as follows; otherwise do not perform step 2 and proceed to step 3 below.

If \(s=28\), the hash value is computed by taking the ones complement of the Exclusive OR of \(\mathrm{VA}_{11: 49}\) with ( \({ }^{11+\mathrm{b}} \mathrm{OIIVA}_{50: 77-\mathrm{b}}\) )
If \(s=40\), the hash value is computed by taking the ones complement of the Exclusive OR of the following three quantities: \(\left(\mathrm{VA}_{24: 37} \mathrm{II}^{25} 0\right)\), ( \(0 \| \mathrm{VA} \mathrm{D}_{0: 37}\) ), and ( \({ }^{b-1} \mathrm{OIIVA}_{38: 77-\mathrm{b}}\) )
The 60-bit real address of a PTEG is formed by concatenating the following values:
■ Bits 0:27 of the 39-bit appropriate primary or secondary hash value ANDed with the mask generated from bits 59:63 of the first doubleword of the Partition Table Entry (HTABSIZE) and then added to the value of bits \(4: 45\) of the first doubleword of the Partition Table Entry (HTABORG).
- Bits 28:38 of the 39 -bit hash value.
- Seven 0-bits.

This operation identifies the "secondary PTEG".
3. As many as 8 PTEs in the primary PTEG and, if the secondary Page Table search is enabled, 8 PTEs in the secondary PTEG are tested to determine if any translate the given virtual address. Let \(q=\) minimum (54, 77-b). For a match to exist, the following conditions must be satisfied, where SLBE is the SLBE used to form the virtual address.
- \(\mathrm{PTE}_{\mathrm{H}}=0\) for the primary PTEG, 1 for the secondary PTEG
- \(\mathrm{PTE}_{\mathrm{V}}=1\)
- PTE \(_{B}=\) SLBE \(_{B}\)
- PTE \(_{\text {AVA }[0: q]}=\mathrm{VA}_{0: q}\)
- if \(b=12\) then
\(\left(P T E_{L}=0\right)\) I ( \(P T E_{L P}\) specifies the \(4 K B\) base page size)
else
\(\left(\mathrm{PTE}_{\mathrm{L}}=1\right)\) \& \(\left(\mathrm{PTE}_{\mathrm{LP}}\right.\) specifies the base page size specified by SLBE \(_{\text {LIILP }}\) )
If no match is found, the search fails. The result is a page fault -- a [Hypervisor] Instruction Storage exception or a [Hypervisor] Data Storage exception, depending on whether the effective address is for an instruction fetch or for a data access. If one
match is found, the search succeeds. If more than one match is found, one of the matching entries is used as if it were the only matching entry, or a Machine Check occurs.

If the Page Table search succeeds, the real address (RA) is formed by concatenating bits 0:59-p of ARPNIILP from the matching PTE with bits 64-p:63 of the effective address (the byte offset), where the \(p\) value is the \(\log _{2}\) (actual page size specified by PTE \(_{\text {LLP }}\) ).
\[
R A=(A R P N \text { II LP })_{0: 59-p} \| E A_{64-p: 63}
\]

A TLB entry may be created as a result of the successful HPT translation. Depending on the specific TLB implementation, the scope of the entry may be the base page size, the virtual page size, or any size in between. In the absence of a TLB, software would be required to create a PTE for each base page sized piece of storage within the virtual page. The number of PTEs actually created to map a virtual page will depend on the scopes supported for TLB entries, the access pattern, and the lifetime of the TLB entries. Hardware generally will not create more than one TLB entry to translate a given virtual address. Multiple matching TLB entries may be created only if the Page Table contains PTEs that map different-sized virtual pages that overlap in the virtual address space. If a TLB search finds multiple matching TLB entries created from such PTEs, one of the matching TLB entries is used as if it were the only matching entry, or a Machine Check occurs. Software should scrupulously avoid creating such mappings.

\section*{Programming Note}

If \(\mathrm{PTE}_{\mathrm{L}}=0\), the actual page size (and base page size) are 4 KB . Otherwise the actual page size and base page size are specified by PTE \(_{\text {LP }}\)
Since hardware searches the Page Table using a value of \(b\) equal to \(\log _{2}\) (base page size specified in the SLBE that was used to translate the address) regardless of the actual page size, the hardware Page Table search will identify different PTEs for VAs in different \(2^{\text {b }}\)-byte blocks of the virtual page if the actual page size is larger than the base page size. Therefore, there may need to be a valid PTE corresponding to each \(2^{\text {b }}\)-byte block of the virtual page that is referenced. For an actual page size that is larger than \(2^{23}(8 \mathrm{MB})\), the PTE \(_{\text {AVA }}\) will differ among some or all of these PTEs. Depending on the Page Table size, some or all of these PTEs may be in the same PTEG. Any such PTEs that are in the same PTEG will differ in the value of PTE \(_{H}\) or PTE \(_{\text {AVA }}\) or both.

All PTEs for the same virtual page should have the same values in the Page Protection, KEY, ARPN, WIMG, and N fields. A set of values from any one of the PTEs that maps the virtual page may be used for an access in the virtual page since lookaside buffer information may be used to translate the virtual address.

To avoid creating multiple matching PTEs, software should not create PTEs for each of two different virtual pages that overlap in the virtual address space. If the virtual page sizes differ, two virtual pages overlap if the values of virtual address bits 0:77-p for both virtual pages are the same, where \(2^{p}\) is the actual virtual page size of the larger page.

\section*{Programming Note}

Because a segment may contain pages of different sizes, the Page Table search uses the segment's base page size (which is the same for all virtual pages in the segment).
■ The value of \(b\) used when searching the Page Table to identify the PTEGs to be checked for a match is \(\log _{2}\) (segment's base page size).
- A PTE (in the selected PTEGs) satisfies the Page Table search only if the base page size specified in the PTE is equal to the segment's base page size.
The matching PTE supplies the actual page size, \(2^{p}\); this value of \(p\) is used in forming the real address.

A virtual page of \(2^{p}\) bytes in a segment with a base page size of \(2^{b}\) bytes may be mapped by as many as \(2^{(p-b)}\) PTEs.

\section*{Programming Note}

To obtain the best performance, Page Table Entries should be allocated beginning with the first empty entry in the primary PTEG, or with the first empty entry in the secondary PTEG if the primary PTEG is full and the secondary Page Table search is enabled ( LPCR \(_{T C}=0\) ).
| In Paravirtualized HPT mode, the N (No-execute) value used for the storage access is the result of ORing the N bit from the matching PTE with the \(N\) bit from the SLB entry that was used to translate the effective address.

\subsection*{5.7.10 Radix Tree Translation}

Radix Tree translation uses a nested set of tables to map storage with increasing granularity. Although there is no requirement for an individual table to have uniform content, Page Directories generally contain pointers to other Page Directories or Page Tables (Page Directory Entries, PDEs), while Page Tables are the leaf tables that contain PTEs. Each Page Directory Entry and Page Table Entry in the Radix Tree is 8 bytes long. A Radix Tree root descriptor (RTRD) specifies the size of the address being translated, the size of the root table, and its location. RTRDs appear in variants of the Partition and Process Table Entries. (See Figures 21 and 22.) The Root Page Directory Size (RPDS) is specified as \(\log _{2}\) (number of entries in the table). That number of bits is taken from the most significant end of the portion of the address being translated, as an index to choose an element in the Root Page Directory. The entries in the Root Page Directory each point to another page of entries, and give its size in the Next Level Size field, \(\mathrm{PDE}_{\text {NLS }}\). The next most significant NLS bits are taken from the address to choose an entry in that table. The process continues until an entry is found that has its Leaf bit set, indicating it is a Page Table Entry. The base size of the page mapped by the PTE is determined by the number of bits remaining in the address after removing the bits used to select the Page Directory and Page Table Entries. An example with RPDS \(=13\) and \(P D E E_{\text {NLS }}=9\) in each Page Directory is shown in Figure 32.
The sizes of table supported at each level of the Radix Tree, as well as the ultimate page sizes supported, are implementation specific with the following exceptions. Implementations must support two Radix Tree configurations that map 52 bit effective addresses: each starting with a 64 KB root page size followed by 2 levels of 4KB tables, ending with either a 256 byte table or a 4 KB table. The former produces a page size of 64 KB and the latter a 4 KB page size. In both cases, a leaf node in the next to last level of table produces a 2MB page size.


Figure 32. Four level Radix Tree walk translating a 52b EA with NLS=13 in the root PDE and NLS=9 in the other PDEs.
5.7.10.1 Radix Tree Page Directory
Entry
\begin{tabular}{|c|c|c|c|c|c|}
\hline V & L & \(/\) & NLB & /// & NLS \\
\hline 0 & 1 & 3 & & 55 & 58 \\
\hline
\end{tabular}
\begin{tabular}{cll} 
Bit(s) & Name & Description \\
0 & V & Valid \\
1 & L & Leaf (entry is a PTE) \\
\(4: 55\) & NLB & Next Level Base \\
\(59: 63\) & NLS & Next Level Size (size of next level \\
& & of table is \(2^{\text {NLS }}+3\) ), NLS \(\geq 5\)
\end{tabular}

All other fields are reserved.
Figure 33. Radix Tree Page Directory Entry

\subsection*{5.7.10.2 Radix Tree Page Table Entry}
\begin{tabular}{|c|c|c|c|cc|c|c|c|c|c|c|}
\hline V & L & sw & // & RPN & sw & R & C & \(/\) & ATT & EAA \\
\hline 0 & 1 & 2 & 6 & & 51 & 54 & 55 & 56 & 57 & 59 & 63
\end{tabular}
\begin{tabular}{|c|c|c|}
\hline Bit(s) & Name & Description \\
\hline 0 & V & Valid \\
\hline 1 & L & Leaf (entry is a PTE) \\
\hline 2 & sw & SW bit 0 \\
\hline 7:51 & RPN & Real Page Number \\
\hline 52:54 & sw & SW bits 1:3 \\
\hline 55 & R & Reference \\
\hline 56 & C & Change \\
\hline 58:59 & Att & \begin{tabular}{l}
Attributes (equivalent WIMG value) \\
Ob00- normal memory (0010) \\
0b01-SAO (1110) \\
Ob10- non-idempotent I/O (0111) \\
Ob11- tolerant I/O (0110)
\end{tabular} \\
\hline 60:63 & EAA & Encoded Access Authority \\
\hline & 0 & \begin{tabular}{l}
Privilege (applies only to pro-cess-scoped translation) \\
0 - problem state access permitted; privileged access if not secure \\
1 - privileged access only
\end{tabular} \\
\hline & 1 & \begin{tabular}{l}
Read 0 - loads not permitted \\
1 - loads permitted
\end{tabular} \\
\hline & 2 & \begin{tabular}{l}
Read/Write \\
0 - loads and stores not permitted \\
1 - loads and stores permitted
\end{tabular} \\
\hline & 3 & \begin{tabular}{l}
Execute \\
0 - instruction execution not permitted 1 - instruction execution permitted
\end{tabular} \\
\hline
\end{tabular}

All other fields are reserved.
Figure 34. Radix Tree Page Table Entry

\subsection*{5.7.11 Nested Translation}

When an operating system that uses Radix Tree translation runs on a hypervisor that uses Radix Tree translation, each guest real address must undergo partition-scoped translation using the hypervisor's Radix Tree for the partition. See Figure 35.


Figure 35. Radix on Radix Page Table search for a 52-bit EA depicting memory reads 1-24 numbered in sequence

When nested translation is being performed, there is the potential for two different sets of protection settings and two different sets of storage attributes. For protection settings, the least permissive values take
effect. For read, write, and execute authority, each is controlled independently based on the least permissive setting of the two translation mechanisms (including all component authority mechanisms within each of them). For storage ordering, the SAO attribute takes effect when both SAO and normal memory attributes are specified. (The hypervisor will typically specify "normal memory" and the OS may override that with SAO.)

When the non-idempotent ( \(\mathrm{IG}=0 \mathrm{~b} 11\) ) and tolerant (IG=0b10) I/O attributes are specified, the setting specified by the guest operating system takes effect. Any other mismatch of attributes will cause a data storage or instruction storage exception, as appropriate for the access.

Reference and Change bit recording is done in both the process-scoped and partition-scoped Page Table Entries. Recording is done as described in Section 5.7.13, "Reference and Change Recording".
Faults that occur as part of process-scoped translation will generally be signalled by ISI/DSI, while faults that occur as part of partition-scoped translation will generally be signalled by HISI/HDSI. (In this categorization, single level translation is considered process-scoped translation except when VPM is active, in which case it is treated like partition-scoped translation.) The address specified in ASDR is the guest real address or VSID for which translation has most immediately failed except when the translation fails too early to produce that value. HDAR will generally contain the EA or lower VA bits for which translation has most immediately failed. For example, in the case of a Page Directory being paged out, the ASDR will contain the guest real address of the Page Directory Entry (down to bit 51), rather than the GRA of the datum being accessed. Exceptions may be manifest in unexpected ways. For example, an instruction fetch can fail to set a Change bit in the host PTE mapping the guest PTE. Similarly, the Reference bit update might fail for lack of write authority on the PTE.

For performance reasons, the result of each walk of a Radix Tree or HPT may be cached in a TLB. Logically, the result of each walk is cached separately. For nested translation, the effective to guest real (pro-cess-scoped) translation may be cached, as well as the partition-scoped translation for each guest real address produced by the translation process. A minimum of two TLB accesses is required to complete a nested translation: one for the effective to guest real address and one for the guest real to host real address. (An implementation may optimize the process, as long as the optimization can be managed correctly using the tlbie instructions that software will use to manage the logical model.)

\subsection*{5.7.12 Translation Process}

As previously described, in its most complicated form the translation process includes the following steps:
- use of the PTCR to find the required Partition Table Entry
- use of the Partition Table Entry to find the parti-tion-scoped Page Table
- use of partition-scoped Page Table to find the required Process Table Entry
- use of the Process Table Entry and parti-tion-scoped Page Table to find the required Seg-
ment Table Entry or walk the process-scoped Page Table (i.e. translate the effective address to a virtual or guest real address), and
- use of the partition-scoped Page Table to translate the virtual or guest real address.

Depending on the translation mode and process state, some of these steps may be skipped. The following subsections enumerate the cases and explain the steps in more detail.

\subsection*{5.7.12.1 Fully-Qualified Address}

The storage control facilities enable hardware to perform the entire translation process given a "fully-qualified address" and context that makes it a unique input. In addition to its normal use, the term "effective address" is sometimes used as shorthand for the fully-qualified address, and the architecture should be read with this possibilty in mind. The following are the components of the fully-qualified address.
- effLPID
- effPID
- EA

The additional context required to perform a translation or match a cached translation may include the following.
- PATE \(_{\text {HR GR }}\) (selected using the value in LPIDR, not effLPID)
- MSR HV PRIR DR \(^{\text {■ }}\)

At a high level, the translation mode is selected by the Host Radix and Guest Radix bits found in the Partition Table Entry. The Host Radix bit indicates whether the hypervisor is using HPT or Radix Tree translation. The Guest Radix bit indicates whether the guest manages an effective to guest real mapping using Radix Tree translation, or an effective to virtual mapping using Segment translation. Given the overall process, \(M_{\text {MVRIPRIIIRIIDR }}\) determine where and how the process is entered.

\subsection*{5.7.12.2 Finding the Page Tables}
[The following description assumes that no legacy mode is active, i.e. LPCR \(_{\text {UPRT }}=1\).]
The components of the fully-qualified address are used to determine the table(s) used in the translation process. The effective LPID and effective PID are used to find the appropriate Page Table base address(es) using the In-Memory Table structures. Some types of translation use process-scoped Page Tables, some use parti-tion-scoped Page Tables, and some use both.
Process-scoped table descriptors are found in the Process Tables as follows. The partition table is assumed to be aligned in host real address space. The Partition Table Entry (PATE) host real address is calculated by ORing the Partition Table Base Address (PATB\|I \({ }^{12} 0\) ) in the PTCR with 16 times the effective LPID and then
performing partition-scoped translation. The second doubleword of the entry contains the base address of the Process Table for the partition. The Process Table is assumed to be aligned in effective ( \(\mathrm{HR}=1\), effLPID=0), virtual, or guest real address space. (If the table is not aligned or is not large enough to support the PID value, an unreported error will most likely result.) The Process Table Entry (PRTE) host real address is calculated by ORing the Process Table Base Address (PRTB \(\|^{40} 0\) for for an HPT host and PRTB \(\|^{12} 0\) for a radix host) in the PATE with 16 times the effective PID and then performing partition-scoped translation. (If the table is not aligned or is not large enough to support the PID value, an unreported error will most likely result.) The resulting Process Table Entry guest real address for a radix guest on a radix host (effLPID \(\neq 0\) ), effective address for radix host (effLPID=0), or virtual address for all cases with an HPT host must be translated via the appropriate partition-scoped table. The Process Table Entry at that location contains a pro-cess-scoped table base address, which is a guest real address for a radix guest on a radix host ( \(\mathrm{HV}=0\) ), a host real address for a radix host ( \(\mathrm{HV}=1\) ), or a virtual address (all cases with an HPT host). The virtual or guest real address must be translated via the appropriate partition-scoped table.

> Programming Note
> The guest real address for a guest of a radix host,or virtual address for a guest of an HPT host, of the Process Table may be set via an hcall. The radix guest may choose to map the Process Table into its own virtual address space. These matters are not visible to the architecture.

\section*{- Programming Note}

Note that the sole purpose of partition-scoped Page Table descriptor when LPID=0 for a radix host is to translate the effective addresses of the Process Table Entries for LPID=0. (If the Process Table Base address for LPID=0 was a real address, the Process Table would have to be in contiguous real storage.) This descriptor will commonly be the same as the descriptor found in the LPID=0, PID=0 Process Table Entry, both pointing to the hypervisor's own page trable, but it may be set up to point to a table used solely to translate the addresses of Process Table Entries.

Partition-scoped Page Table descriptors are found in the Partition Table as follows. The Partition Table Base Address is found in the PTCR. The effective LPID (times 16 bytes per partition) is used to index off the Partition Table Base Address to find the appropriate Partition Table Entry. The first doubleword of the entry contains the base address of the Page Table.

\subsection*{5.7.12.3 Obtaining Host Real Address, Radix on Radix}

The following cases exist.
■ Guest access to quadrant 0 with translation on: process-scoped translation is performed on LPIDRIIPIDRIIEA, with the result subject to parti-tion-scoped translation with effective LPID=LPIDR.
- Guest access to quadrant 3 with translation on: process-scoped translation is performed on LPIDRIIOIIEA, with the result subject to parti-tion-scoped translation with effective LPID=LPIDR.
- Hypervisor access to quadrant 1 with translation on: process-scoped translation is performed on LPIDRIIPIDRIIEA, with the result subject to parti-tion-scoped translation with effective LPID=LPIDR if LPIDR \(=0\).
- Hypervisor access to quadrant 2 with translation on: process-scoped translation is performed on LPIDRIIOIIEA, with the result subject to parti-tion-scoped translation with effective LPID=LPIDR if LPIDR \(\neq 0\).
■ Guest OS access with translation off: parti-tion-scoped translation is performed with effective LPID = LPIDR.
- Hypervisor or host application access to quadrant 0 with translation on: process-scoped translation is performed on OlIPIDRIIEA.
- Hypervisor or host application access to quadrant 3 with translation on: process-scoped translation is performed with OlIOIIEA.
- Hypervisor real mode access: subject to HRMOR and \(E A_{0}\) as described in Section 5.7.3.1.


Figure 36. Radix on Radix translation, general case

\subsection*{5.7.12.4 Obtaining Host Real Address, HPT}

There are two scenarios for Paravirtualized HPT translation. The first is the legacy scenario with a native HPT hypervisor. The second scenario is for a Radix Tree translation hypervisor providing a Paravirtualized HPT environment for the guest. In this latter scenario, the LPID=0 Partition Table Entry will have \(\mathrm{HR}=1\) and \(\mathrm{GR}=1\). For both scenarios when \(\mathrm{MSR}_{\mathrm{HV}}=1\), the LPID value is always taken from LPIDR and the PID value is always taken from PIDR. In the latter scenario, the hypervisor will explicitly set LPIDR=0 when it wants to use its Radix Tree(s).
When using Paravirtualized HPT translation, the pro-cess-scoped Page Tables are replaced by Segment Tables, and the description in Section 5.7.12.2, "Finding the Page Tables" can be read with that substitution in mind. The process-scoped translation is the effec-tive-to-virtual translation described in Section 5.7.8. In-Memory Table walks are processed via the LPID=LPIDR partition-scoped HPT.
As with the previous enumerations, this is done from a hardware point of view. As a result, it does not differentiate the software cases for which Segment translation should only be satisfied by bolted translations
The following cases exist.
■ Guest access with translation on: process-scoped translation is performed on LPIDRIIPIDRIIEA with
the result subject to partition-scoped translation using parameters from the matching segment descriptor.
- Hypervisor or adjunct access with translation on and LPID \(\neq 0\) : process-scoped translation, limited to an SLB search with no Segment Table walk, is performed on LPIDRIIPIDRIIEA, with the result subject to partition-scoped translation using parameters from the matching segment descriptor.
- Hypervisor or adjunct access with translation on and LPID=0: process-scoped translation (with Segment Table walk) is performed on LPIDRIIPIDRIIEA, with the result subject to parti-tion-scoped translation using parameters from the matching segment descriptor.
■ Guest OS access with translation off: subject to VPM, as described in Section 5.7.3.3.
■ Hypervisor real mode access: subject to HRMOR and \(E A_{0}\) as described in Section 5.7.3.1.


Figure 37. Paravirtualized HPT translation

\subsection*{5.7.13 Reference and Change Recording}
| When operating in Paravirtualized HPT mode, Reference (R) and Change (C) bits are updated in any one of what could be multiple (because of the multiple base size PTEs mapping a virtual page) Page Table Entries that map the virtual page that is being accessed. When operating in Radix on Radix mode, Reference (R) and Change (C) bits may be updated in multiple Page Table

Entries that are accessed as part of the translation process. (For example, each access to a guest's Page Directory or Page Table Entry potentially sets a Reference bit in the partition-scoped table mapping it.) If the storage operand of a Load or Store instruction crosses a virtual page boundary, the accesses to the components of the operand in each page are treated as separate and independent accesses to each of the pages for the purpose of setting the Reference and Change bits.

For Radix Tree translation, the Reference and Change bits are set atomically, as though the PTE was read to perform the translation using a Load And Reserve instruction, and conditional on the translation being valid and correct (and on the existence of the reservation), the appropriate bit(s) are set as though with a Store Conditional instruction. For HPT translation, Reference and Change bits are set as though the PTE was read to perform the translation using a (simple) Load instruction and the appropriate bit(s) are set as though with a (simple) Store instruction. These accesses may be for a byte, halfword, word, doubleword, or quadword (the same size for both operations of a pair). Setting the bits need not be atomic with respect to performing the access that causes the bits to be updated. The Reference bit must contain 1 in order to load from the corresponding page. The Change bit must contain 1 in order to store to the corresponding page. If hardware is unable to set the bit(s) atomically for Radix Tree translation, a [Hypervisor] Data Storage or [Hypervisor] Instruction Storage interrupt will be caused.

\section*{Programming Note}

The interrupt indicates to software that it must set the appropriate bit(s) itself. Note that an instruction fetch can cause a Change bit to be set, for example in the host Page Table Entry that maps the guest Page Table Entry if the instruction fetch causes the Reference bit to be set in the guest Page Table Entry.

> Programming Note
> The atomic setting of the Reference and Change bits enables an optimized sampling of them, for example when determining what pages to reclaim for other uses. To accurately sample the bits under HPT translation, it is necessary to first invalidate the PTE and the corresponding TLB entries. The optimized sequence eliminates the requirement for the relatively expensive invalidation of the TLB entries before sampling the bits. Instead, software may simply load the PTE using a Load And Reserve instruction, and then set the PTE invalid using a Store Conditional instruction. The TLB invalidation may be defered indefinitely and grouped into clusters or range bombs for improved performance. The Reference and Change bits sampled in this manner are accurate (if the store conditional succeeds) because with the PTE marked invalid, it will be impossible to access a page for which the appropriate bit is not already set.

\section*{- Programming Note}

In nested Radix Tree translation, as many as three Change bits may be set: in the process-scoped and partition-scoped PTEs for the access itself, and in the partition-scoped PTE that maps the pro-cess-scoped PTE. Similarly, a large number of Reference bits may be set, including for each parti-tion-scoped PTE that maps a process-scoped PDE or PTE.

Reference and Change bits are set by the hardware as I described below. An attempt to access storage may cause one or more of the bits to be set (as described below) even if the access is not performed. The bits are updated in the Page Table Entry if the new value would otherwise be different from the old value for the virtual page, as determined by examining either the Page Table Entry or any lookaside information for the virtual page (e.g., TLB) maintained by the hardware.

\section*{Reference Bit}

The Reference bit is set to 1 if the corresponding access (load, store, implicit access, or instruction fetch) is required by the sequential execution model and is performed. Otherwise the Reference bit may be set to 1 if the corresponding access is attempted, either in-order or out-of-order, even if the attempt causes an exception, except that the Reference bit is not set to 1 for the access caused by an indexed Move Assist instruction for which the XER specifies a length of zero.

\section*{Change Bit}

The Change bit is set to 1 if a Store instruction is executed and the store is performed or if an implicit update is performed. Otherwise in general the Change bit may be set to 1 if a Store instruction is executed and the store is permitted by the storage protection mechanism and, if the Store instruction is executed out-of-order, the instruction would be required by the sequential execution model in the absence of the following kinds of interrupts:
■ system-caused interrupts (see Section 6.4 on page 1055)
■ Floating-Point Enabled Exception type Program interrupts when the thread is in an Imprecise mode.
The only exception to the preceding statement is that the Change bit is not set to 1 if the instruction is a Store String Indexed instruction for which the XER specifies a length of zero.

\section*{Programming Note}

A virtual page in a segment with a smaller base page size may be mapped by multiple PTEs. For each access of a virtual page, hardware may search the Page Table to update the \(R\) and \(C\) bits. If lookaside buffer information for the virtual page already indicates that all such bits to be set have already been set in a PTE that maps the virtual page, hardware need not make an update. Consider the following sequence of events:
1. A virtual page is mapped by 2 PTEs A and B and the \(R\) and \(C\) bits in both PTEs are 0 .
2.A Load instruction accesses the virtual page and the \(R\) bit is updated in PTE A.
3.A Load instruction accesses the virtual page and the \(R\) bit is updated in PTE B.
4.A Store instruction accesses the virtual page and the C bit is updated in PTE B .
5. The virtual page is paged out. Software must examine both PTE A and B to get the state of the \(R\) and \(C\) bits for the virtual page.
Furthermore, if in event 2, PTE A was not found, a Data Storage interrupt or Hypervisor Data Storage interrupt may occur. Subsequently, if in event 3 or 4, PTE B was not found, a Data Storage interrupt or Hypervisor Data Storage interrupt may occur.

\section*{Programming Note}

Even though the execution of a Store instruction causes the Change bit to be set to 1, the store might not be performed or might be only partially performed in cases such as the following.
- A Store Conditional instruction (stwcx. or stdcx.) is executed, but no store is performed.
- The Store instruction causes a Data Storage exception (for which setting the Change bit is not prohibited).
- The Store instruction causes an Alignment exception.
- The Page Table Entry that translates the virtual address of the storage operand is altered such that the new contents of the Page Table Entry preclude performing the store (e.g., the PTE is made invalid, or the PP bits are changed).

For example, when executing a Store instruction, the thread may search the Page Table for the purpose of setting the Change bit and then re-execute the instruction. When reexecuting the instruction, the thread may search the Page Table a second time. If the Page Table Entry has meanwhile been altered, by a program executing on another thread, the second search may obtain the new contents, which may preclude the store.
- A system-caused interrupt occurs before the store has been performed.

When the hardware updates the Reference and Change bits in the Page Table Entry, the accesses are performed as described in Section 5.7.3.4, "Storage Control Attributes for Implicit Storage Accesses" on
I page 986. These Reference and Change bit updates are not necessarily immediately visible to software. Executing a sync instruction ensures that all Reference and Change bit updates associated with address translations that were performed, by the thread executing the sync instruction, before the sync instruction is executed will be performed with respect to that thread before the sync instruction's memory barrier is created. There are additional requirements for synchronizing Reference and Change bit updates in multi-threaded systems; see Section 5.10, "Translation Table Update Synchronization Requirements" on page 1043.

\section*{Programming Note}

Because the sync instruction is execution synchronizing, the set of Reference and Change bit updates that are performed with respect to the thread executing the sync instruction before the memory barrier is created includes all Reference and Change bit updates associated with instructions preceding the sync instruction.

\section*{Version 3.0}

If software refers to a Page Table Entry when \(\mathrm{MSR}_{\mathrm{DR}}=1\), the Reference and Change bits in the associated Page Table Entry are set as for ordinary loads and stores. See Section 5.10 for the rules software must follow when updating Reference and Change bits.

Figure 38 on page 1010 summarizes the rules for setting the Reference and Change bits. The table applies to each atomic storage reference. It should be read from the top down; the first line matching a given situation applies. For example, if stwcx. fails due to both a storage protection violation and the lack of a reservation, the Change bit is not altered.

In the figure, the "Load-type" instructions are the Load instructions described in Books I, II, and III, and the Cache Management instructions that are treated as Loads. The "Store-type" instructions are the Store instructions described in Books I, II, and III, and the Cache Management instructions that are treated as Stores. The "ordinary" Load and Store instructions are those described in Books I, II, and III. "set" means "set to 1 ".


Figure 38. Setting the Reference and Change bits

\subsection*{5.7.14 Storage Protection}

The storage protection mechanism provides a means for selectively granting instruction fetch access, granting read access, granting write access, and prohibiting access to areas of storage based on a number of control criteria.
The operation of the storage protection mechanism depends on the contents of one or more of the following.
- MSR bits HV, IR, DR, PR
- the key bits in the associated SLB entry
- the page protection bits and key bits in the associated PTE
I
- the AMR, IAMR, AMOR, and UAMOR

The storage protection mechanism consists of the Virtual Page Class Key Protection mechanism described

\section*{I} in Section 5.7.14.1, the Basic Storage Protection mechanism described in Section 5.7.14.2 and Section 5.7.14.3, and the Radix Tree Translation Storage Protection mechanism described in Section 5.7.14.4.

When address translation is enabled for an access, the
| access is permitted in Paravirtualized HPT mode if and only if the access is permitted by both the Virtual Page Class Key Protection mechanism and the Basic Storage Protection mechanism. When address translation is enabled for a guest access, the access is permitted in Radix on Radix mode if and only if the access is permitted by the Radix Tree Translation Storage Protection mechanism for both the process-scoped and parti-tion-scoped PTEs. When address translation is disabled for a guest access or is enabled for an access with \(\mathrm{MSR}_{\mathrm{HV}}=1\), the access is permitted in Radix on Radix mode if and only if the access is permitted by the Radix Tree Translation Storage Protection mechanism for the partition-scoped PTE. When address translation is disabled for an access with \(M S R_{H V}=1\), the access is permitted if and only if the access is permitted by the Basic Storage Protection mechanism. If an instruction fetch is not permitted, an Instruction Storage exception or a Hypervisor Instruction Storage exception is generated. If a data access is not permitted, a Data Storage exception or a Hypervisor Data Storage exception is generated.
A protection domain is a maximal range of effective addresses for which variables related to storage protection can be independently specified (including by default, as in real and hypervisor real addressing modes), or a maximal range of addresses, effective or virtual, for which variables related to storage protection cannot be specified. Examples include: a segment, a virtual page (including for a virtualized Real Mode Area), the Real Mode Area (regardless of whether the RMA is virtualized), the effective address range \(0: 2^{60}-1\) in hypervisor real addressing mode, and a maximal range of effective or virtual addresses that cannot be
mapped to real addresses. A protection boundary is a boundary between protection domains.

\subsection*{5.7.14.1 Virtual Page Class Key Protection}

The Virtual Page Class Key Protection mechanism provides the means to assign virtual pages to one of 32 classes, and to modify data access permissions for each class by modifying the Authority Mask Register (AMR), shown in Figure 39, and to modify instruction access permissions for each class by modifying the Instruction Authority Mask Register (IAMR) shown in Figure 40.

\section*{Programming Note}

If address translation is disabled for a given access, the access is not affected by the Virtual Page Class Key Protection mechanism even if the access is made in virtual real addressing mode.

\section*{Authority Mask Register}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline Key0 & Key1 & Key2 & & & Key29 & Key30 & Key31 \\
\hline 0 & 2 & 4 & 6 & & 58 & 60 & 62 \\
\hline
\end{tabular}
\begin{tabular}{lll} 
Bits & Name & Description \\
\(0: 1\) & Key0 & Access mask for class number 0 \\
\(2: 3\) & Key1 & Access mask for class number 1 \\
\(\ldots\) & \(\ldots\) & \(\ldots\) \\
\(2 n: 2 n+1\) & Keyn & Access mask for class number \(n\) \\
\(\ldots\) & \(\ldots\) & \(\ldots\) \\
\(62: 63\) & Key31 & Access mask for class number 31
\end{tabular}

\section*{Figure 39. Authority Mask Register (AMR)}

The access mask for each class defines the access permissions that apply to loads and stores for which the virtual address is translated using a Page Table Entry that contains a Key field value equal to the class number. The access permissions associated with each class are defined as follows, where \(\mathrm{AMR}_{2 n}\) and \(A M R_{2 n+1}\) refer to the first and second bits of the access mask corresponding to class number \(n\).
- A store is permitted if \(\mathrm{AMR}_{2 \mathrm{n}}=0 \mathrm{bO}\); otherwise the store is not permitted.
- A load is permitted if \(\mathrm{AMR}_{2 n+1}=0 \mathrm{~b} 0\); otherwise the load is not permitted.
The AMR can be accessed using either SPR 13 or SPR 29. Access to the AMR using SPR 29 is privileged.

\section*{Programming Note}

Because the AMR is part of the program context (if address translation is enabled), and because it is desirable for most application programmers not to have to understand the software synchronization requirements for context alterations (or the nuances of address translation and storage protection), operating systems should provide a system library program that application programs can use to modify the AMR.

\section*{Instruction Authority Mask Register}
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline Key0 & Key1 & Key2 & & \(\ldots\) & Key29 & Key30
\end{tabular} Key31 \begin{tabular}{l} 
K.
\end{tabular}
\begin{tabular}{lll} 
Bits & Name & Description \\
\(0: 1\) & Key0 & Access mask for class number 0 \\
\(2: 3\) & Key1 & Access mask for class number 1 \\
\(\ldots\) & \(\ldots\) & \(\ldots\) \\
\(2 n: 2 n+1\) & Keyn & Access mask for class number \(n\) \\
\(\ldots\) & \(\ldots\) & \(\ldots\) \\
\(62: 63\) & Key31 & Access mask for class number 31
\end{tabular}

Figure 40. Instruction Authority Mask Register (IAMR)

The access mask for each class defines the access permissions that apply to instruction fetches for which the virtual address is translated using a Page Table Entry that contains a Key field value equal to the class number. The access permission associated with each class is defined as follows, where \(\mathrm{IAMR}_{2 n+1}\) refers to the bit of the access mask corresponding to class number n .
- An instruction fetch is permitted if \(\mathrm{IAMR}_{2 \mathrm{n}+1}=\mathrm{ObO}\); otherwise the instruction fetch is not permitted.

Bit 0 of each key field is reserved
Access to the IAMR is privileged.

The Authority Mask Override Register (AMOR) and the User Authority Mask Override Register (UAMOR), shown in Figure 41 and Figure 42 respectively, can be used to restrict modifications (mtspr) of the AMR. Also, the AMOR can be used to restrict modifications of the UAMOR and IAMR. Access to both the AMOR and UAMOR is privileged. The AMOR is a hypervisor resource.


Figure 41. Authority Mask Override Register (AMOR)
\begin{tabular}{|ll|}
\hline \multicolumn{3}{|c|}{ UAMOR } \\
\hline 0 & 63
\end{tabular}

Figure 42. User Authority Mask Override Register (UAMOR)

The bits of the AMOR and UAMOR are in 1-1 correspondence with the bits of the AMR (i.e., [U]AMOR \({ }_{i}\) corresponds to \(\mathrm{AMR}_{\mathrm{i}}\) ). The AMOR affects modifications of the AMR and UAMOR in privileged but non hypervisor state; the UAMOR affects modifications of the AMR in problem state.

Similarly, the odd bits of the AMOR are in 1-1 correspondence with the odd bits of the IAMR (i.e., \(\mathrm{AMOR}_{2 j+1}\) corresponds to \(I \mathrm{AMR}_{2 j+1}\) ). The AMOR affects modifications of the IAMR in privileged but non hypervisor state; the IAMR cannot be accessed in problem state.
■ When mtspr specifying the AMR (using either SPR 13 or SPR 29) or the IAMR is executed in privileged but non-hypervisor state, the AMOR is used as a mask that controls which bits of the resulting AMR or IAMR contents come from register RS and which AMR or IAMR bits are not modified.
- Similarly, when mtspr specifying the AMR (using SPR 13) is executed in problem state, the UAMOR is used as a mask that controls which bits of the resulting AMR contents come from register RS and which AMR bits are not modified.
- When mtspr specifying the UAMOR is executed in privileged but non-hypervisor state, the AMOR is ANDed with the contents of register RS and the result is placed into the UAMOR; the AMOR thereby controls which bits of the resulting UAMOR contents come from register RS and which UAMOR bits are set to zero.
A complete description of these effects can be found in the description of the mtspr instruction in Section 4.4.5.

Software must ensure that both bits of each even/odd bit pair of the AMOR contain the same value. - i.e., the contents of register RS for mtspr specifying the AMOR must be such that \((R S)_{2 n}=(R S)_{2 n+1}\) for every \(n\) in the range 0:31 - and likewise for the UAMOR. If this
requirement is violated for the UAMOR the results of accessing the UAMOR (including implicitly by the hardware as described in the second item of the preceding list) are boundedly undefined; if the requirement is violated for the AMOR the results of accessing the AMOR (including implicitly by the hardware as described in the first and third items of the list) are undefined.

\section*{Programming Note}

The preceding requirement permits designs to implement the AMOR and/or UAMOR as 32-bit registers - specifically, to implement only the even-numbered bits (or only the odd-numbered bits) of the register - in a manner such that the reduction, from the architecturally-required 64 bits to 32 bits, is not visible to (correct) software. This implementation technique saves space in the hardware. (A design that uses this technique does the appropriate "fan in/out" when the register is accessed, to provide the appearance, to (correct) software, of supporting all 64 bits of the register.)

Permitting designs to implement the [U]AMOR as 32-bit registers by virtue of the software requirement specified above, rather than by defining the [U]AMOR as 32-bit registers, permits the architecture to be extended in the future to support controlling modification of the "read access" AMR bits (the odd-numbered bits) independently from the "write access" AMR bits (the even-numbered bits), if that proves desirable. If this independent control does prove desirable, the only architecture change would be to eliminate the software requirement.

\section*{Programming Note}

When modifying the AMOR and/or UAMOR, the hypervisor should ensure that the two registers are consistent with one another before giving control to a non-hypervisor program. In particular, the hypervisor should ensure that if \(\mathrm{AMOR}_{\mathrm{i}}=0\) then UAMOR \(_{i}=0\), for all \(i\) in the range \(0: 63\). (Having \(A M O R_{i}=0\) and \(U A M O R_{i}=1\) would permit problem state programs, but not the operating system, to modify AMR bit i.)

\section*{Programming Note}

The Virtual Page Class Key Protection mechanism replaces the Data Address Compare mechanism that was defined in versions of the architecture that precede Version 2.04 (e.g., the two facilities use some of the same resources, as described below). However, the Virtual Page Class Key Protection mechanism can be used to emulate the Data Address Compare mechanism. Moreover, programs that use the Data Address Compare mechanism can be modified in a manner such that they will work correctly both on implementations that comply with versions of the architecture that precede Version 2.04 (and hence implement the Data Address Compare mechanism) and on implementations that comply with Version 2.04 of the architecture or with any subsequent version (and hence instead implement the Virtual Page Class Key Protection mechanism). The technique takes advantage of the facts that the SPR number for privileged access to the AMR (29) is the same as the SPR number for the Data Address Compare mechanism's ACCR (Address Compare Control Register), that KEY \({ }_{4}\) occupies the same bit in the PTE as the Data Address Compare mechanism's AC (Address Compare) bit, and that the definition of \(A^{\prime} C R_{62: 63}\) is very similar to the definition of each even-odd pair of AMR bits. The technique is as follows, where PTE1 refers to doubleword 1 of the PTE.
- Set bits 2:3 and 62:63 of SPR 29 (which is either the ACCR or the AMR) to \(x\), where \(x\) is the desired 2 -bit value for controlling Data Address Compare matches, and set bits 0:1 to Os.
- Set PTE1 \({ }_{54}\) (which is either the AC bit or \(\mathrm{KEY}_{4}\) ) to the same value that the AC bit would be set to, and set \(\mathrm{PTE}_{2: 3}\) (which are either RPN bits, that correspond to a real address size larger than the size supported by any implementation that supports the Data Address Compare mechanism, or \(\mathrm{KEY}_{0: 1}\) ) and PTE \(1_{52: 53}\) (which are either reserved bits or \(\mathrm{KEY}_{2: 3}\) ) to 0 s .
- Use PTE KEY values 0 and 1 only for purposes of emulating the Data Address Compare mechanism, except that PTE \(_{\text {KEY }}\) value 0 may
also be used for any virtual pages for which it is desired that the Virtual Page Class Key Protection mechanism permit all accesses. Do not use PTE \(_{\text {KEY }}=31\).
- When a Hypervisor Data Storage interrupt occurs, if \(\mathrm{HDSISR}_{42}=1\) then ignore the interrupt for Cache Management instructions other than dcbz. (These instructions can cause a virtual page class key protection violation but cannot cause a Data Address Compare match.) Otherwise forward the interrupt to the operating system, which will treat the interrupt as if a Data Address Compare match had occurred. (Note: Cases for which it is undefined whether a Data Address Compare match occurs do not necessarily cause a virtual page class key protection violation.)
(Because privileged software can access the AMR using either SPR 13 or SPR 29, it might seem that, when SPR 13 was added to the architecture (in Version 2.06), SPR 29 should have been removed. SPR 29 is retained for two reasons: first, to avoid requiring privileged software to change to use the newer SPR number; and second, to retain the ability to emulate the Data Address Compare mechanism as described above.)

\begin{abstract}
\section*{Programming Note}

An example of the use of the AMOR (and UAMOR) is to support adjuncts (see Section 5.7.4, "Definitions"), The hypervisor could use KEY value j for all data virtual pages that only the adjunct must be able to access. Before dispatching the partition for the first time, the hypervisor would initialize the three registers as follows.
AMR: all 0 s except bits 2 j and \(2 \mathrm{j}+1\), which would contain 1s
UAMOR: all Os
AMOR: all 1 s except bits 2 j and \(2 \mathrm{j}+1\), which would contain \(0 s\)

Before dispatching the adjunct, the hypervisor would set UAMOR to all 0s, and would set the AMR to all 1 s except bits 2 j and \(2 \mathrm{j}+1\), which would be set the AMOR and the adiunct cannot modify the UAMOR.) In addition, the hypervisor would prevent the partition from modifying or deleting PTEs that contain translations used by the adjunct.
(It may be desirable to avoid using KEY values 0,1 , and 31 for storage that only the adjunct can access, because these KEY values may be needed by the partition to emulate the Data Address Compare mechanism, as described above. Also, old software, that was written for an implementation that complies with a version of the architecture that precedes Version 2.04 (the version in which virtual page class keys were added), effectively uses KEY 0 for all virtual pages.)
\end{abstract}

\section*{Programming Note}

Initialization of the UAMOR to all 0s, by the hypervisor before dispatching a partition for the first time, as described in the preceding Programming Note, permits operating systems (in partitions that run in a compatibility mode corresponding to Version 2.06 of the architecture or a subsequent version) to migrate gradually to supporting problem state access to the AMR - specifically, to avoid having to be changed immediately to modify the UAMOR and to save the AMR contents when an interrupt occurs from problem state. Relatedly, having the UAMOR contain all Os while an application program is running protects old application programs that are "AMR-unaware". In the absence of programming errors, such application programs would not attempt to read or modify the AMR. However, having the UAMOR contain all Os protects such programs against modifying the AMR inadvertently.
Permitting an "AMR-unaware" application program to modify the AMR (inadvertently) is potentially harmful for the obvious reasons. (The program might set to 1 an AMR bit corresponding to accesses that are necessary in order for the program to work correctly.) Moreover, even for an operating system that includes support for problem state modification of the AMR, having the UAMOR contain all Os allows the operating system to avoid saving and restoring the AMR for "AMR-unaware" application programs. Such an operating system would provide a system service program that allows an application program to declare itself to be "AMR-aware" - i.e., potentially to need to modify the AMR. When an application program invokes this service, the operating system would set the UAMOR to the non-zero value appropriate to the access authorities (load and/or store, for one or more key values) that the application program is allowed to modify, and thereafter would save and restore the AMR (and preserve the UAMOR) for this application program. (Having the UAMOR contain all Os does not prevent an "AMR-unaware" program from reading the AMR, but inadvertent reading of the AMR is likely to be much less harmful than inadvertently modifying it.)
(For partitions that run in a compatibility mode corresponding to a version of the architecture that precedes Version 2.06, the PCR provides sufficient protection to application programs.)

\subsection*{5.7.14.2 Basic Storage Protection, Address Translation Enabled}

When address translation is enabled, , the Basic Storage Protection mechanism is controlled by the following.
- MSR \(_{\text {PR }}\), which distinguishes between supervisor (privileged) state and problem state
- \(\mathrm{K}_{\mathrm{s}}\) and \(\mathrm{K}_{\mathrm{p}}\), the supervisor (privileged) state and problem state storage key bits in the SLB entry used to translate the effective address
- PP, page protection bits 0:2 in the Page Table Entry used to translate the effective address
- For instruction fetches only:
- the N (No-execute) value used for the access (see Sections 5.7.8.1 and 5.7.9.2)
- \(\mathrm{PTE}_{\mathrm{G}}\), the G (Guarded) bit in the Page Table Entry used to translate the effective address

Using the above values, the following rules are applied.
1. For an instruction fetch, the access is not permitted if the \(N\) value is 1 or if \(\mathrm{PTE}_{\mathrm{G}}=1\).
2. For any access except an instruction fetch that is not permitted by rule 1, a "Key" value is computed using the following formula:
\[
\text { Key } \leftarrow\left(K_{p} \& M S R P R\right) \mid\left(K_{s} \& \neg M S R_{P R}\right)
\]

Using the computed Key, Figure 43 is applied. An instruction fetch is permitted for any entry in the figure except "no access". A load is permitted for any entry except "no access". A store is permitted only for entries with "read/write".
\begin{tabular}{|c|c|l|}
\hline Key & PP & Access Authority \\
\hline 0 & 000 & read/write \\
0 & 001 & read/write \\
0 & 010 & read/write \\
0 & 011 & read only \\
0 & 110 & read only \\
\hline 1 & 000 & no access \\
1 & 001 & read only \\
1 & 010 & read/write \\
1 & 011 & read only \\
1 & 110 & no access \\
\hline
\end{tabular}

All PP encodings not shown above are reserved. The results of using reserved PP encodings are boundedly undefined.

Figure 43. PP bit protection states, address translation enabled

\subsection*{5.7.14.3 Basic Storage Protection, Address Translation Disabled}

When address translation is disabled, the Basic Storage Protection mechanism is controlled by \(\mathrm{MSR}_{\mathrm{HV}}\), which (when MSR \({ }_{P R}=0\) ) distinguishes between hypervisor state and privileged but non-hypervisor state (see Chapter 2 and Section 5.7.3, "Hypervisor Real And Vir-
| tual Real Addressing Modes"). The following rules apply.
1. If \(\mathrm{MSR}_{H V}=0\), access authority is determined as described in Section 5.7.3.3.
2. If \(M S R_{H V}=1\), the access is permitted.

\subsection*{5.7.14.4 Radix Tree Translation Storage Protection}

The storage protection mechanism for Radix Tree translation is completely different from what is provided for HPT translation. EAA \(1: 3\) provide control over read, read/write, and execute access if the process has the appropriate privilege. \(E A A_{0}\), together with the security setting in the AMR or IAMR, provide three protection configurations for process-scoped translation: (1) insecure mode gives equivalent access to privileged and problem state processes, (2) secure mode gives access only to problem state, and (3) access only to privileged processes. (Note that privileged includes hypervisor privileged.) For partition-scoped translation, including translation of table entry addresses, either value of \(E A A_{0}\) permits the access. See Figure 34 and Figure 44 for details. The selection between 1 and DC for process-scoped protection of privileged read and write is determined by key 0 of the AMR. When bit 0 is 0 , the privileged bit in the PTE is ignored for a privileged store. When bit 0 is 1 , the privileged bit must be 1 for a privileged store. Similarly when bit 1 is 0 , the privileged bit in the PTE is ignored for a privileged load. When bit 1 is 1 , the privileged bit must be 1 for a privileged load. The selection between 1 and DC for pro-cess-scoped protection of execute is determined by key 0 of the IAMR. When bit 1 is 0 , the privileged bit in the PTE is ignored for an attempt to execute the instruction in privileged state. When bit 1 is 1 , the privileged bit must be 1 to execute the instruction in priivleged state.


Figure 44. Encoded Access Authority (aka page protection)

\subsection*{5.7.15 Cluster Shared Memory Protection}

Any access to Cluster Shared Memory by instructions other than copy and paste[.] is prohibited. Instruction fetches and implicit storage accesses must not address CSM. Access to CSM by copy andpaste[.] is subject to the storage protection mechanisms described in Section 5.7.14. The platform may provide an additional authorization mechanism for these accesses. If it does, the access is permitted if and only if it is permitted by both kinds of mechanism.

\section*{Programming Note}

CSM is to be marked No-execute by the hypervisor, so that an instruction fetch will violate storage protection rather than having boundedly undefined behavior.

\subsection*{5.8 Storage Control Attributes}

This section describes aspects of the storage control attributes that are relevant only to privileged software programmers. The rest of the description of storage control attributes may be found in Section 1.6 of Book II and subsections.

\subsection*{5.8.1 Guarded Storage}

Storage is said to be "well-behaved" if the corresponding real storage exists and is not defective, and if the effects of a single access to it are indistinguishable from the effects of multiple identical accesses to it. Data and instructions can be fetched out-of-order from well-behaved storage without causing undesired side effects.

Storage is said to be Guarded if any of the following conditions is satisfied.
- MSR bit IR or DR is 1 for instruction fetches or data accesses respectively, and the \(G\) bit is 1 in the relevant HPT Page Table Entry.
- MSR bit IR or DR is 1 for instruction fetches or data accesses respectively, and the Att field has the value 0b10 in the relevant Radix Page Table Entry.
■ MSR bit IR or DR is 0 for instruction fetches or data accesses respectively, \(\mathrm{MSR}_{\mathrm{HV}}=1\), and the storage is outside the range(s) specified by the Hypervisor

Real Mode Storage Control facility (see Section 5.7.3.2.1).

In general, storage that is not well-behaved should be Guarded. Because such storage may represent a control register on an I/O device or may include locations that do not exist, an out-of-order access to such storage may cause an I/O device to perform unintended operations or may result in a Machine Check.

The following rules apply to in-order execution of Load and Store instructions for which the first byte of the storage operand is in storage that is both Caching Inhibited and Guarded.
- Load or Store instruction that causes an atomic access

If any portion of the storage operand has been accessed and an External, Decrementer, Hypervisor Decrementer, Performance Monitor, or Imprecise mode Floating-Point Enabled exception is pending, the instruction completes before the interrupt occurs.
- Load or Store instruction that causes an Alignment exception, or that causes a [Hypervisor] Data Storage exception for reasons other than Data Address Watchpoint match.

The portion of the storage operand that is in Caching Inhibited and Guarded storage is not accessed.
(The corresponding rules for instructions that cause a Data Address Watchpoint match are given in Section 8.4.)

\subsection*{5.8.1.1 Out-of-Order Accesses to Guarded Storage}

In general, Guarded storage is not accessed out-of-order. The only exceptions to this rule are the following.

\section*{Load Instruction}

If a copy of any byte of the storage operand is in a cache then that byte may be accessed in the cache or in main storage.

\section*{Instruction Fetch}

If \(\mathrm{MSR}_{H V} \mathrm{IR}^{=}=0 \mathrm{~b} 10\) then an instruction may be fetched if any of the following conditions are met.
1. The instruction is in a cache. In this case it may be fetched from the cache or from main storage.
2. The instruction is in a real page from which an instruction has previously been fetched, except that if that previous fetch was based on condition 1 then the previously fetched instruction must have been in the instruction cache.
3. The instruction is in the same real page as an instruction that is required by the sequential execu-
tion model, or is in the real page immediately following such a page.

\begin{abstract}
\section*{Programming Note}

Software should ensure that only well-behaved storage is copied into a cache, either by accessing as Caching Inhibited (and Guarded) all storage that may not be well-behaved, or by accessing such storage as not Caching Inhibited (but Guarded) and referring only to cache blocks that are well-behaved.

If a real page contains instructions that will be executed when \(\mathrm{MSR}_{\mathrm{IR}^{\prime}}=0\) and \(M S R_{H V}=1\), software should ensure that this real page and the next real page contain only well-behaved storage (or that the Hypervisor Real Mode Storage Control facility specifies that this real page is not Guarded).
\end{abstract}

I

\subsection*{5.8.2 Storage Control Bits}

When address translation is enabled, each storage access is performed under the control of the Page Table Entry used to translate the effective address. Each Page Table Entry contains storage control bits that specify the presence or absence of the corresponding storage control for all accesses translated by the entry as shown in Figure 45 and Figure 46. In the following description, references to individual WIMG bits apply to the corresponding Radix Att encoding except where otherwise stated or obvious from context.
\begin{tabular}{|l|l|}
\hline Bit & Storage Control Attribute \\
\hline \(\mathrm{W}^{1,3}\) & \begin{tabular}{l}
\(0-\) not Write Through Required \\
\(1-\) Write Through Required
\end{tabular} \\
\hline \(\mathrm{I}^{3}\) & \begin{tabular}{l}
\(0-\) not Caching Inhibited \\
\(1-\) - Caching Inhibited
\end{tabular} \\
\hline \(\mathrm{M}^{2}\) & \begin{tabular}{l}
\(0-\) not Memory Coherence Required \\
\(1-\) Memory Coherence Required
\end{tabular} \\
\hline G & \begin{tabular}{l}
\(0-\) not Guarded \\
\(1-\) Guarded
\end{tabular} \\
\hline 1 \begin{tabular}{l} 
Support for the 1 value of the W bit is optional. \\
Implementations that do not support the 1 value \\
treat the bit as reserved and assume its value to \\
be 0. \\
Support for the 0 value of the M bit is optional, \\
implementations that do not support the 0 value \\
assume the value of the bit to be 1, and may either \\
preserve the value of the bit or write it as 1. \\
The combination WIMG = Ob1110 has behavior \\
unrelated to the meanings of the individual bits. \\
See see Section 5.8.2.1, "Storage Control Bit \\
Restrictions" for additional information.
\end{tabular} \\
\hline
\end{tabular}

I Figure 45. Storage control bits, HPT PTE
\begin{tabular}{|l|lc|}
\hline Att value & Storage Type & \\
\hline 00 & normal memory & \((\mathrm{WIMG}=0010)\) \\
\hline \(01^{1}\) & SAO & \((\mathrm{WIMG}=1110)\) \\
\hline 10 & non-idempotent I/O \(\quad(\mathrm{WIMG}=0111)\) \\
\hline 11 & tolerant I/O & \((\mathrm{WIMG}=0110)\) \\
\hline \(\mathrm{W}=0\) always for Radix Tree translation \\
\(\mathrm{M}=1\) always for Radix Tree translation \\
1 & \multicolumn{1}{l|}{\begin{tabular}{l} 
Behaves \\
order.
\end{tabular}} & \\
\hline
\end{tabular}

Figure 46. Storage control bits, Radix PTE
When address translation is enabled, instructions are not fetched from storage for which the \(G\) bit in the Page Table Entry is set to 1; see Section 5.7.14.

When address translation is disabled, the storage control attributes are implicit; see Section 5.7.3.2.
In Sections 5.8.2.1 and 5.8.2.2, "access" includes accesses that are performed out-of-order, and references to W, I, M, and G bits include the values of those bits that are implied when address translation is disabled.

\section*{Programming Note}

In a system consisting of only a single-threaded processor which has caches, correct coherent execution does not require storage to be accessed as Memory Coherence Required, and accessing storage as not Memory Coherence Required may give better performance.

\subsection*{5.8.2.1 Storage Control Bit Restrictions}

All combinations of \(\mathrm{W}, \mathrm{I}, \mathrm{M}\), and \(G\) values are permitted except those for which both \(W\) and \(I\) are 1 and \(M I I G \neq 0 b 10\).
The combination WIMG \(=0 \mathrm{~b} 1110\) is used to identify the Strong Access Ordering (SAO) storage attribute (see Section 1.7.1, "Storage Access Ordering", in Book II). Because this attribute is not intended for general purpose programming, it is provided only for a single combination of the attributes normally identified using the WIMG bits. That combination would normally be indicated by WIMG = 0b0010.
References to Caching Inhibited storage (or storage with \(\mathrm{I}=1\) ) elsewhere in the Power ISA have no application to SAO storage or its WIMG encoding, despite the encoding using \(\mathrm{I}=1\). Conversely, references to storage that is not Caching Inhibited (or storage with \(\mathrm{I}=0\) ) apply to SAO storage or its WIMG encoding. References to Write Through Required storage (or storage with \(\mathrm{W}=1\) ) elsewhere in the Power ISA have no application to SAO storage or its WIMG encoding, despite the fact that the
encoding uses \(\mathrm{W}=1\). Conversely, references to storage that is not Write Through Required (or storage with \(\mathrm{W}=0\) ) apply to SAO storage or its WIMG encoding.
If a given real page is accessed concurrently as SAO storage and as non-SAO storage, the result may be characteristic of the weakly consistent model.

\section*{Programming Note}

If an application program requests both the Write Through Required and the Caching Inhibited attributes for a given storage location, the operating system should set the I bit to 1 and the W bit to 0 . The operating system should provide a means by which application programs can request SAO storage, in order to avoid confusion with the preceding guideline (since SAO is encoded using WI=0b11).

At any given time, the value of the W bit must be the same for all accesses to a given real page.
At any given time, the value of the I bit must be the same for all accesses to a given real page.

\subsection*{5.8.2.2 Altering the Storage Control Bits}

When changing the value of the W bit for a given real page from 0 to 1 , software must ensure that no thread modifies any location in the page until after all copies of locations in the page that are considered to be modified in the data caches have been copied to main storage using dcbst or dcbf[I].
When changing the value of the I bit for a given real page from 0 to 1 , software must set the I bit to 1 and then flush all copies of locations in the page from the caches using dcbf[I] and icbi before permitting any other accesses to the page. Note that similar cache management is required before using the Fixed-Point Load and Store Caching Inhibited instructions to access storage that has formerly been cached. (See Section 4.4.1 on page 965.)

\section*{Programming Note}

The storage control bit alterations described above are examples of cases in which the directives for application of statements about the W and I bits to SAO given in the third paragraph of the preceding subsection must be applied. A transition from the typical WIMG=0b0010 for ordinary storage to WIMG=0b1110 for SAO storage does not require the flush described above because both WIMG combinations indicate storage that is not Caching Inhibited.

\section*{Version 3.0}

\section*{Programming Note}

It is recommended that dcbf be used, rather than dcbfl, when changing the value of the I or W bit from 0 to 1. (dcbfl would have to be executed on all threads for which the contents of the data cache may be inconsistent with the new value of the bit, whereas, if the M bit for the page is 1, dcbf need be executed on only one thread in the system.)

When changing the value of the M bit for a given real page, software must ensure that all data caches are consistent with main storage. The actions required to do this are system-dependent.

\section*{Programming Note}

For example, when changing the \(M\) bit in some directory-based systems, software may be required to execute dcbf[] on each thread to flush all storage locations accessed with the old \(M\) value before permitting the locations to be accessed with the new \(M\) value.

Additional requirements for changing the storage control bits in the Page Table are given in Section 5.10.

\subsection*{5.9 Storage Control Instructions}

\subsection*{5.9.1 Cache Management Instructions}

This section describes aspects of cache management that are relevant only to privileged software programmers.

For a dcbz instruction that causes the target block to be newly established in the data cache without being fetched from main storage, the hardware need not verify that the associated real address is valid. The existence of a data cache block that is associated with an invalid real address (see Section 5.6) can cause a
delayed Machine Check interrupt or a delayed Checkstop.

Each implementation provides an efficient means by which software can ensure that all blocks that are considered to be modified in the data cache have been copied to main storage before the thread enters any power conserving mode in which data cache contents are not maintained.

\subsection*{5.9.2 Synchronize Instruction}

The Synchronize instruction is described in Section 4.6.3 of Book II, but only at the level required by an application programmer. This section describes properties of the instruction that are relevant only to operating system and hypervisor software programmers.

The Synchronize instruction provides an ordering function for stores that are in set A of the memory barrier created by the Synchronize instruction, relative to data accesses caused by instructions that are executed on other threads after the occurrence of the interrupt that is caused by a msgsndp or msgsnd instruction that follows the Synchronize instruction. The thread that is the target of the msgsndp or msgsnd instruction is here called the "target thread".
- For msgsndp, and \(L=0,1\), or 2 for the Synchronize instruction, the stores are performed with respect to the target thread before any data accesses caused by instructions that are executed on the target thread after the corresponding Directed Privileged Doorbell interrupt has occurred.
- For msgsnd, and \(L=0\) or 2 for the Synchronize instruction (sync or ptesync), the stores are performed with respect to any given other thread before any data accesses caused by instructions that are executed on the given thread after a msgsync instruction is executed on that thread after the corresponding Directed Hypervisor Doorbell interrupt has occurred on the target thread.

\section*{Programming Note}

Synchronize with L=1 (Iwsync) should not be used with msgsnd. (If used, it will not have the desired ordering effect.)

\section*{- Programming Note}

The msgsync instruction, which is needed when msgsnd is used, is not needed when msgsndp is used because msgsndp targets only threads on the same multi-threaded processor as the thread executing the msgsndp, while msgsnd can target any thread in the system. (If the target thread for msgsnd is on the same multi-threaded processor as the thread executing the msgsnd, in principle the msgsync can be omitted. This optimization is practical only when the msgsnd topology is appropriately constrained, however, because the Directed Hypervisor Doorbell interrupt provides no indication of which thread executed the msgsnd that caused the interrupt, so there is no easy way for the interrupt handler to determine whether the msgsync can be omitted.) msgsync is not needed or defined in V. 2.07 for a similar reason: msgsnd in V. 2.07 can target only threads on the same multi-threaded processor as the thread executing the msgsnd.
The ordering done by sync (and ptesync) provides the appearance of "causality" across a sequence of msgsnd instructions, as in the following example. "msgsnd->T1" means "msgsnd instruction targetting thread T1". "<DHDI 0>" means "occurrence of Directed Hypervisor Doorbell interrupt caused by msgsnd executed on TO". On TO, register r1 is assumed to contain the value 1.
\begin{tabular}{lll} 
T0 & T1 & T2 \\
std r1,X & <DHDI 0> & <DHDI 1> \\
sync & msgsnd->T2 & msgsync \\
msgsnd->T1 & & ld r1,X
\end{tabular}

In this example, T2's load from X must return 1.
Another variant of the Synchronize instruction is described below. It is designated the Page Table Entry Synchronize instruction, and is specified by the
extended mnemonic ptesync (equivalent to sync with \(\mathrm{L}=2\) ).
The ptesync instruction has all of the properties of sync with \(\mathrm{L}=0\) and also the following additional properties.
■ The memory barrier created by the ptesync instruction provides an ordering function for the storage accesses associated with all instructions that are executed by the thread executing the ptesync instruction and, as elements of set A, for all Reference and Change bit updates associated with additional address translations that were performed, by the thread executing the ptesync instruction, before the ptesync instruction is executed. The applicable pairs are all pairs \(\mathrm{a}_{\mathrm{i}}, \mathrm{b}_{\mathrm{j}}\) in which \(b_{j}\) is a data access and \(a_{i}\) is not an instruction fetch.

■ The ptesync instruction causes all Reference and Change bit updates associated with address translations that were performed, by the thread executing the ptesync instruction, before the ptesync instruction is executed, to be performed with respect to that thread before the ptesync instruction's memory barrier is created.
The memory barrier created by the ptesync instruction provides an ordering function for all stores to the Partition Table, Process Tables, Segment Tables, Page Directories, and Page Tables caused by Store instructions preceding the ptesync instruction with respect to invalidations, of cached copies of information derived from these tables, caused by slbieg and tlbie instuctions following the ptesync instruction. The memory barrier ensures that all searches of these tables by another thread, that are performed after an invalidation caused by such an slbieg or tlbie instruction has been performed with respect to the other thread and that implicitly load from the target location of such a store, will obtain the value stored (or a value stored subsequently).

\section*{- Programming Note}

The next bullet is sufficient to order the stores with respect to the invalidations on the thread executing the ptesync instruction. That bullet is also sufficient to provide the ordering with respect to invalidations caused by slbie, slbia, and tlbiel instuctions, which affect only the thread executing them.
- The ptesync instruction provides an ordering function for all stores to the Partition Table, Process Tables, Segment Tables, Page Directories, and Page Tables caused by Store instructions preceding the ptesync instruction with respect to searches of these tables that are performed, by the thread executing the ptesync instruction, after the ptesync instruction completes. Executing a pte-
sync instruction ensures that all such searches that implicitly load from the target location of such a store will obtain the value stored (or a value stored subsequently). Also, the memory barrier created by the ptesync insruction ensures that all searches of these tables by any other thread, that are performed after a store in set B of the memory barrier has been performed with respect to the other thread and that implicitly load from the target location of such a store, will obtain the value stored (or a value stored subsequently).
■ In conjunction with the tlbie and tlbsync instructions, the ptesync instruction provides an ordering function for TLB invalidations and related storage accesses on other threads as described in the tlbsync instruction description on page 1042.

Similarly, in conjunction with the slbieg and slbsync instructions, the ptesync instruction provides an ordering function for SLB invalidations and related storage accesses on other threads as described in the slbsync instruction description on page 1031.

\section*{Programming Note}

For instructions following a ptesync instruction, the memory barrier need not order implicit storage accesses for purposes of address translation and reference and change recording.
The functions performed by the ptesync instruction may take a significant amount of time to complete, so this form of the instruction should be used only if the functions listed above are needed. Otherwise sync with \(L=0\) should be used (or sync with \(L=1\), or eieio, if appropriate).

Section 5.10, "Translation Table Update Synchronization Requirements" on page 1043 gives examples of uses of ptesync.

\subsection*{5.9.3 Lookaside Buffer Management}

All implementations have a Segment Lookaside Buffer (SLB). Independent of whether the executing partition operates in a mode that uses hardware SLB loading and bolting versus pure software loading (controlled by the value of LPCR \(_{\text {UPRT }}\) ), software is responsible for keeping the SLB current with the segment mapping for the process that is executing. To simplify management of the SLB, hardware will only create speculative SLB entries for the context that is executing, in particular using the current values of LPID and PID. Except when LPID \(=0\), no such entries will be created when \(\mathrm{MSR}_{\mathrm{HV}}=1\). Similarly, when \(\mathrm{MSR}_{\mathrm{IR}}=0\) and \(\mathrm{MSR}_{\mathrm{DR}}=0\), no such entries will be created. Proper management of the SLB across context switches is described in programming notes.

For performance reasons, most implementations also cache other information that is used in address translation. These caches may include: a Translation Lookaside Buffer (TLB) which is a cache of recently used Page Table Entries (PTEs); a cache of recently used translations of effective addresses to real addresses; a Page Walk Cache for Radix Tree translation; caching of the In-Memory Tables; or any combination of these. Lookaside information, including the SLB, is managed using the instructions described in the subsections of this section unless additional requirements are provided in implementation-specific documentation.

Lookaside information derived from PTEs is not necessarily kept consistent with the Page Table. When software alters the contents of a PTE, in general it must
| also invalidate all corresponding TLB entries and imple-mentation-specific lookaside information; exceptions to this rule are described in Section 5.10.1.2.
| The effects of the slbie, slbieg, slbia, and TLB Management instructions on address translations, as specified in Sections 5.9.3.2 for the SLB and 5.9.3.3 for the TLB, Page Walk Cache, and In-Memory Table caches, apply to all implementation-specific lookaside information that is used in address translation. Unless otherwise stated or obvious from context, references to SLB entry invalidation and TLB entry invalidation elsewhere in the Books apply also to invalidation of Page Walk Cache content, In-Memory Table cache content, and all implementation-specific lookaside information that is derived from SLB entries and PTEs, respectively.
| All implementations provide a means by which software can invalidate all implementation-specific lookaside information that is derived from PTEs.

Implementation-specific lookaside information that contains translations of effective addresses to real addresses may include "translations" that apply in real addressing mode. Because such "translations" are affected by the contents of the LPCR and HRMOR, when software alters the (relevant) contents of these registers it must also invalidate the corresponding implementation-specific lookaside information. Software can invalidate all such lookaside information by using the slbia instruction with \(\mathrm{IH}=0 \mathrm{~b} 000\). However, performance is likely to be better if other, appropriate, IH values are used to limit the amount of lookaside information that invalidated.

All implementations that have such lookaside information provide a means by which software can invalidate all such lookaside information.

For simplicity, elsewhere in the Books it is assumed that the TLB exists.

\section*{Programming Note}

Because the instructions used to manage TLBs, SLBs, Page Walk Caches, caches of Partition and Process Table Entries, and implementation-specific lookaside information may be changed in a future version of the architecture, it is recommended that software "encapsulate" their use into subroutines.

\section*{Programming Note}

The function of all the instructions described in Sections 5.9.3.2 - 5.9.3.3 is independent of whether address translation is enabled or disabled.

For a discussion of software synchronization requirements when invalidating SLB and TLB entries, see Chapter 11.

\subsection*{5.9.3.1 Thread-Specific Segment Translations}

It is necessary to provide thread-specific temporary ESID to VSID translations. These translations cannot be placed in valid entries in the Segment Table because the Segment Table has a process scope rather than a thread scope. Instead, software will use slbmte to install such translations in the SLB. All SLB entries created using slbmte are considered to be "software created." Software created entries will only translate accesses from the hardware thread by which they are installed. When LPCR \({ }_{\text {UPRT }}=1\), they are also considered to be "bolted." Each thread has the ability to bolt four entries.

\subsection*{5.9.3.2 SLB Management Instructions}

Software establishes translations in the SLB using slbmte. Care must be taken to avoid creating multiple effective-to-virtual translations for any given effective address. Software-created entries will remain in the SLB until invalidated using slbie or slbia (which also invalidate related implementation-specific lookaside information) or overwritten using slbmte. After updating a Segment Table Entry, software must use an slbie or slbieg instruction to remove lookaside information associated with the old contents of the entry. slbie may be used to invalidate software-created entries, but will not invalidate outboard translation caches. slbieg does not invalidate software-created entries, but is the only way to invalidate outboard translation caches. slbsync will establish order between slbieg instructions and a subsequent ptesync. ptesync must also be used to synchronize the Segment Table update prior to performing the lookaside management. When performing a context switch, software must use an slbia instruction to remove lookaside information associated with the old context. slbmfee and slbmfev may be used by the hypervisor to save software-created entries. slbmte is used to restore software-created
entries. slbfee has no function when \(\operatorname{LPCR}_{\text {UPRT }}=1\) for the partition that is running.

\section*{Programming Note}

Accesses to a given SLB entry caused by the instructions described in this section obey the sequential execution model with respect to the contents of the entry and with respect to data dependencies on those contents. That is, if an instruction sequence contains two or more of these instructions, when the sequence has completed, the final contents of the SLB entry and of General Purpose Registers is as if the instructions had been executed in program order.

However, software synchronization is required in order to ensure that any alterations of the entry take effect correctly with respect to address translation; see Chapter 11.

\section*{Programming Note}

Changes to the segment mappings in the presence of active transactions may compromise transactional semantics if the transaction has accessed a segment that is assigned a new VSID. Consequently, when modifying segment mappings, it is the responsibility of the OS or hypervisor to ensure that any transaction that may have touched the modified segment is terminated, using a tabort. or treclaim. instruction.

SLB Invalidate Entry
X-form
slbie RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{1}{|l|}{ I/I } & & I/I & RB & & 434 \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
```

ea 0:35}\leftarrow(\textrm{RB}\mp@subsup{)}{0:35}{
if, for SLB entry that translates
or most recently translated ea,
entry_class = (RB)36 and
entry_seg_size = size specified in (RB) 37:38
then for SLB entry (if any) that translates ea
SLBE
all other fields of SLBE \leftarrow undefined
else
s \leftarrow log_base_2(entry_seg_size)
esid}\leftarrow(\textrm{RB}\mp@subsup{)}{0:63-s}{-
u}\leftarrow\mathrm{ undefined 1-bit value
if u then
if an SLB entry translates esid
SLBE
all other fields of SLBE }\leftarrow\mathrm{ undefined

```

Let the Effective Address (EA) be any EA for which \(E A_{0: 35}=(R B)_{0: 35}\). Let the class be \((R B)_{36}\). Let the segment size be equal to the segment size specified in \((R B)_{37: 38}\); the allowed values of \((R B)_{37: 38}\), and the correspondence between the values and the segment size, are the same as for the B field in the SLBE (see Figure 26 on page 995).

The class value and segment size must be the same as the class value and segment size in the SLB entry that translates the EA, or the values that were in the SLB entry that most recently translated the EA if the translation is no longer in the SLB; if these values are not the same, it is implementation-dependent whether the SLB entry (or implementation-dependent translation information) that translates the EA is invalidated, and the next paragraph need not apply.
If the SLB contains only a single entry that translates the EA, then that is the only SLB entry that is invalidated, except that it is implementation-dependent whether an implementation-specific lookaside entry for a real mode address "translation" is invalidated. If the SLB contains more than one such entry, then zero or more such entries are invalidated, and similarly for any implementation-specific lookaside information used in address translation; additionally, a machine check may occur.

SLB entries are invalidated by setting the V bit in the entry to 0 , and the remaining fields of the entry are set to undefined values.

This instruction terminates any Segment Table walks being performed on behalf of the thread that executes it.
The hardware ignores the contents of RB listed below and software must set them to 0s.
- \(\quad(\mathrm{RB})_{37}\)
- \(\quad(\mathrm{RB})_{39}\)
- \(\quad(\mathrm{RB})_{40: 63}\)
- If \(s=40,(R B)_{24: 35}\)

If this instruction is executed in 32-bit mode, \((R B)_{0: 31}\) must be zeros.

This instruction is privileged.

\section*{Special Registers Altered:}

None
Programming Note
slbie does not affect SLBs on other threads.

\section*{Programming Note}

The reason the class value specified by slbie must be the same as the Class value that is or was in the relevant SLB entry is that the hardware may use these values to optimize invalidation of implemen-tation-specific lookaside information used in address translation. If the value specified by slbie differs from the value that is or was in the relevant SLB entry, these optimizations may produce incorrect results. (An example of implementation-specific address translation lookaside information is the set of recently used translations of effective addresses to real addresses that some implementations maintain in an Effective to Real Address Translation (ERAT) lookaside buffer.)
When switching tasks in certain cases, it may be advantageous to preserve some implementa-tion-specific lookaside entries while invalidating others. The \(\mathrm{IH}=0 \mathrm{~b} 001\) invalidation hint of the slbia instruction can be used for this purpose if SLB class values are appropriately assigned, i.e., a class value of 0 gives the hint that the entry should be preserved and a class value of 1 indicates the entry must be invalidated. Also, it is advantageous to assign a class value of 1 to entries that need to be invalidated via an slbie instruction while preserving implementation-specific lookaside entries that are not derived from an SLB entry since such entries are assigned a class value of 0 .

\section*{Programming Note}

The B value in register RB may be needed for invalidating ERAT entries corresponding to the translation being invalidated.

\section*{Programming Note}

When switching to execute an adjunct, a hypervisor will turn off translation and use slbie to be sure there is no SLB entry mapping the effective address space that will be used by the incoming adjunct. It will then bolt an entry for the incoming adjunct and transfer control to that adjunct. While translation is off and during adjunct execution, no speculative Segment Table walks will be performed.

\section*{SLB Invalidate Entry Global \\ X-form \\ slbieg RS,RB \\ \begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & \multicolumn{2}{|c|}{ RS } & \multicolumn{1}{|l|}{} & RB & & 466 \\
\hline 0 & & & 11 & 16 & 21 & \\
31 \\
\hline
\end{tabular}}
```

target_PID $=\mathrm{RS}_{0: 31}$
if $\mathrm{MSR}_{\mathrm{HV}}=1$ then target_LPID $=\mathrm{RS}_{32}: 63$
else target_LPID = LPIDR
$\mathrm{ea}_{0: 35} \leftarrow(\mathrm{RB})_{0: 35}$
for each thread with LPIDR=target_LPID and
PIDR=target_PID
if, for each SLB entry that
translates or most recently translated ea
entry_class $=(R B) 36$ and
entry_seg_size = size specified in (RB) $37: 38$
then for SLB entry (if any)
that translates ea and is not software-created
$\mathrm{SLBE}_{\mathrm{V}} \leftarrow 0$
all other fields of SLBE $\leftarrow$ undefined
else
$\mathrm{s} \leftarrow \log _{\text {_base_2 }}$ (entry_seg_size)
esid $\leftarrow(R B)_{0: 63-s}$
$u \leftarrow$ undefined 1-bit value
if $u$ then
if an SLB entry translates esid and the entry
is not software-created
$\mathrm{SLBE}_{\mathrm{V}} \leftarrow 0$
all other fields of SLBE $\leftarrow$ undefined

```

The operation performed by this instruction is based on the contents of registers RS and RB. The contents of these registers are shown below.

RS
\begin{tabular}{|l|ll|}
\hline PID & \multicolumn{1}{c|}{ LPID } \\
\hline 0 & 32 & 63 \\
\hline
\end{tabular}

RB
\begin{tabular}{ll|l|l|l|l|}
\hline & ESID & C & B & Os \\
\hline 0 & & 3637 & 39 & & 63 \\
RS \(_{0: 31}\) & PID & & & & \\
\hline
\end{tabular}
\(\mathrm{RS}_{32: 63}\) LPID
\(\mathrm{RB}_{0: 35}\) ESID
\(\mathrm{RB}_{36} \quad \mathrm{C}\)
\(\mathrm{RB}_{37: 38} \mathrm{~B}\)
\(\mathrm{RB}_{39: 63}\) must be 0b0 II 0x000000

Let the target PID be \(\mathrm{RS}_{0: 31}\). If the instruction is executed in hypervisor state, let the target LPID be \(\mathrm{RS}_{32: 63}\); otherwise let the target LPID be the contents of LPIDR. Let the Effective Address (EA) be any EA for which \(E A_{0: 35}=(R B)_{0: 35}\). Let the class be \((R B)_{36}\). Let the segment size be equal to the segment size specified in (RB) 37:38 ; the allowed values of (RB) 37:38 , and the correspondence between the values and the seg-
ment size, are the same as for the B field in the SLBE (see Figure 26 on page 995).
Only SLBs for threads running on behalf of target_LPID and target_PID are searched. Software-created entries are ignored. The class value and segment size must be the same as the class value and segment size in the SLB entry that translates the EA, or the values that were in the SLB entry that most recently translated the EA if the translation is no longer in the SLB; if these values are not the same, it is implementation-dependent whether the SLB entry (or implementation-dependent translation information) that translates the EA is invalidated, and the next paragraph need not apply.
If the SLB contains only a single entry that translates the EA, then that is the only SLB entry that is invalidated, except that it is implementation-dependent whether an implementation-specific lookaside entry for a real mode address "translation" is invalidated. If the SLB contains more than one such entry, then zero or more such entries are invalidated, and similarly for any implementation-specific lookaside information used in address translation; additionally, a machine check may occur.

SLB entries are invalidated by setting the V bit in the entry to 0 , and the remaining fields of the entry are set to undefined values.

The hardware ignores the contents of RB listed below and software must set them to 0 s .
\[
\begin{array}{ll}
- & (\mathrm{RB})_{37} \\
- & (\mathrm{RB})_{39} \\
- & (\mathrm{RB})_{40: 63} \\
- & \text { If } s=40,(\mathrm{RB})_{24: 35}
\end{array}
\]

If this instruction is executed in 32-bit mode, \((\mathrm{RB})_{0: 31}\) must be zeros.
This instruction is privileged except when \(\mathrm{LPCR}_{\text {GTSE }}=0\), making it hypervisor privileged.

\section*{Special Registers Altered:}

None

\section*{Programming Note}
slbieg does affect SLBs on other threads.

\begin{abstract}
Programming Note
The reason the class value specified by slbieg must be the same as the Class value that is or was in the relevant SLB entry is that the hardware may use these values to optimize invalidation of imple-mentation-specific lookaside information used in address translation. If the value specified by slbieg differs from the value that is or was in the relevant SLB entry, these optimizations may produce incorrect results. (An example of implementation-specific address translation lookaside information is the set of recently used translations of effective addresses to real addresses that some implementations maintain in an Effective to Real Address Translation (ERAT) lookaside buffer.)
When switching tasks in certain cases, it may be advantageous to preserve some implementa-tion-specific lookaside entries while invalidating others. The IH=0b001 invalidation hint of the slbia instruction can be used for this purpose if SLB class values are appropriately assigned, i.e. a class value of 0 gives the hint that the entry should be preserved and a class value of 1 indicates the entry must be invalidated. Also, it is advantageous to assign a class value of 1 to entries that need to be invalidated via an slbieg instruction while preserving implementation-specific lookaside entries that are not derived from an SLB entry since such entries are assigned a class value of 0 .
\end{abstract}

\section*{Programming Note}
slbieg specifying LPID \(=0\) and \(\mathrm{PID}=0\) in RS should be used to invalidate hypervisor real mode ERAT entries in the "nest" (see the appropriate platform architecture document), since hypervisor real mode ERAT entries in the processor are typically invalidate using slbie, which does not broadcast.

\section*{Programming Note}

The \(B\) value in register RB may be needed for invalidating ERAT entries corresponding to the translation being invalidated.

\section*{Programming Note}

Use of slbieg to invalidate software-created segment descriptors is a programming error. The architecture requires that bolted entries be ignored by the instruction.

\section*{SLB Invalidate AII}

X-form
slbia IH
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 31 & I/ & IH & I/I & I/I & 498 & \(/\) \\
\hline
\end{tabular}
\begin{tabular}{|l|l|l|l|l|l|l|}
\hline 0 & 6 & 8 & 11 & 16 & 21 & 31 \\
\hline
\end{tabular}
```

for each SLB entry except SLB entry 0
SLBE
all other fields of SLBE \leftarrow undefined

```

For all SLB entries except SLB entry 0 , the V bit in the entry is set to 0 , making the entry invalid, and the remaining fields of the entry are set to undefined values. SLB entry 0 is altered only if \(\mathrm{IH}=0 \mathrm{~b} 100\) or its C bit is 1 and \(\mathrm{IH}=0 \mathrm{~b} 011\).

On implementations that have implementation-specific lookaside information for effective to real address translations, the IH field provides a hint that can be used to invalidate entries selectively in such lookaside information. The defined values for IH are as follows.

Ob000 All such implementation-specific lookaside information is invalidated. (This value is not a hint.)

Ob001 Preserve such implementation-specific lookaside information having a Class value of 0 .
Ob010 Preserve such implementation-specific lookaside information created when \(\mathrm{MSR}_{\mathrm{IR} / D \mathrm{R}}=0\).

Ob011 Preserve such implementation-specific lookaside information having a Class value of 0 . Preserve SLB entries with a Class value of 0 . If SLB entry 0 has a Class value of 1 , it is invalidated. (This value is not a hint.)
Ob100 All such implementation-specific lookaside information is invalidated, and in addition, SLB entry 0 is invalidated. (This value is not a hint.)
Ob110 Preserve such implementation-specific lookaside information created when \(\mathrm{MSR}_{\mathrm{HV}}=1\), \(M S R_{P R}=0\), and \(M S R_{I R / D R}=0\).
Ob111 All such implementation-specific lookaside information is invalidated, but no SLB entries are invalidated.

All other IH values are reserved. If the IH field contains a reserved value, the hint provided by the IH field is undefined.

Implementation specific lookaside information for which preservation is not requested is invalidated. Implementation specific lookaside information for which preservation is requested may be invalidated.

When \(\mathrm{IH}=0 \mathrm{~b} 000\), execution of this instruction has the side effect of clearing the storage access history associated with the Hypervisor Real Mode Storage Control facility. See Section 5.7.3.2.1, "Hypervisor Real Mode Storage Control" for more details.
This instruction terminates any Segment Table walks being performed on behalf of the thread that executes it, and ensures that any new table walks will be performed using the current PIDR value.

\section*{Programming Note}

When performing a context switch between processes, an operating system will use mtPIDR followed by slbia. The synchronization of the PID value and termination of outstanding Segment Table walks ensures that SLB will not contain multiple entries mapping the same EA range (i.e. from the former and new PIDs). Note that if this sequence is performed with translation enabled, care must be taken to avoid an implicit branch. (i.e. the same translation(s) for the locations containing the context switch routine must be valid for both processes.)

For the corresponding situation when changing partitions, hypervisor software should get all the affected threads into real mode, execute mtLPIDR, and then perform the slbia on all the affected threads. Avoiding the implicit branch is too difficult.

This instruction is privileged.

\section*{Special Registers Altered:}

None

\section*{Programming Note}
slbia does not affect SLBs on other threads.

\section*{I a Programming Note}

If \(\boldsymbol{s}\) lbia is executed when instruction address translation is enabled, software can ensure that attempting to fetch the instruction following the slbia does not cause an Instruction Segment interrupt by placing the slbia and the subsequent instruction in the effective segment mapped by SLB entry 0 . (The preceding assumes that no other interrupts occur between executing the slbia and executing the subsequent instruction. It also assumes that IH values other than 0b011 and 0b100 are used.)

\section*{Programming Note}

The defined values for IH are as follows.
Ob000 All ERAT entries are invalidated. (This value is not a hint.)

0b001 Preserve ERAT entries with a Class value of 0 . This value should be used by an operating system with backward compatibility requirements when switching tasks in certain cases; for example, if \(\mathrm{SLBE}_{\mathrm{C}}=0\) is used for SLB translations shared between the tasks.

Ob010 Preserve ERAT entries created when \(\mathrm{MSR}_{\mathrm{IR} / \mathrm{DR}}=0\). This value should generally be used by an operating system when switching tasks.

Ob011 Preserve ERAT and SLB entries with a Class value of 0 . If SLB entry 0 has \(\mathrm{SLBE}_{\mathrm{C}}=1\), it is invalidated. (This value is not a hint.) This value should be used by an operating system without backward compatibility requirements when switching tasks in certain cases; for example, if \(\mathrm{SLBE}_{\mathrm{C}}=0\) is used for SLB translations shared between the tasks.

Ob100 All ERAT entries are invalidated, and in addition, SLB entry 0 is invalidated. (This value is not a hint.) This value should be used by the hypervisor when relocating itself (i.e. when modifying the HRMOR) or when reconfiguring real storage.

Ob110 Preserve ERAT entries created when \(\mathrm{MSR}_{\mathrm{HV}}=1\) and \(\mathrm{MSR}_{\mathrm{IR} / \mathrm{DR}}=0\). This value should be used by the hypervisor when switching partitions.

Ob111 All ERAT entries are invalidated. No SLB entries are invalidated. (This setting is provided mainly for use prior to product shipment.)

\section*{Programming Note}
slbia serves as both a basic and an extended mnemonic. The Assembler will recognize an slbia mnemonic with one operand as the basic form, and an slbia mnemonic with no operand as the extended form. In the extended form the IH operand is omitted and assumed to be 0 .
SLB Move To Entry X-form
slbmte RS,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RS & \multicolumn{1}{|l|}{ I/I } & RB & & 402 & \\
\hline 0 & & 6 & 11 & 16 & 21 & \\
\hline
\end{tabular}

When LPCR \({ }_{\text {UPRT }}=0\), this instruction is the sole means for specifying Segment translations to the hardware. When LPCR \({ }_{\text {UPRT }}=1\), Segment Table walks populate the SLB, and this instruction is used only to bolt thread-specific Segment translations.
The SLB entry specified by bits \(52: 63\) of register RB is loaded from register RS and from the remainder of register RB. The contents of these registers are interpreted as shown in Figure 47.
RS
\begin{tabular}{|l|ll|l|l|l|ll|}
\hline B & & VSID & \(\mathrm{K}_{\mathrm{s}} \mathrm{K}_{\mathrm{p}}\) NLC & 0 & LP & \multicolumn{2}{|c|}{ Os } \\
\hline 0 & 2 & & 52 & 57 & 58 & 60 & 63 \\
\hline
\end{tabular}

RB
\begin{tabular}{|l|l|l|ll|}
\hline ESID & V & Os & \multicolumn{2}{c|}{ index } \\
\hline 0 & 3637 & 52 & 63 \\
\hline
\end{tabular}
\(\mathrm{RS}_{0: 1} \quad \mathrm{~B}\)
\(\mathrm{RS}_{2: 51}\) VSID
\(R_{52} \quad K_{s}\)
\(\begin{array}{ll}R_{53} & \mathrm{~K}_{\mathrm{p}} \\ \mathrm{RS}_{54} & \mathrm{~N}\end{array}\)
\(\mathrm{RS}_{55} \quad \mathrm{~L}\)
\(R S_{56} \quad C\)
\(\mathrm{RS}_{57}\) must be 0b0
\(\mathrm{RS}_{58: 59} \mathrm{LP}\)
\(\mathrm{RS}_{60: 63}\) must be 0b0000
\(\mathrm{RB}_{0: 35}\) ESID
\(\mathrm{RB}_{36} \mathrm{~V}\)
\(\mathrm{RB}_{37: 51}\) must be \(0 b 000\) II \(0 x 000\)
\(\mathrm{RB}_{52: 63}\) index, which selects the SLB entry
Figure 47. GPR contents for slbmte
On implementations that support a virtual address size of only \(n\) bits, \(n<78,(R S)_{2: 79-n}\) must be zeros.
When LPCR \(_{\text {UPRT }}=1\), the value of index must not exceed 3. (RB) 52:61 are ignored.
High-order bits of \((\mathrm{RB})_{52: 63}\) that correspond to SLB entries beyond the size of the SLB provided by the implementation must be zeros.

The hardware ignores the contents of RS and RB listed below and software must set them to 0 s.
\(\begin{array}{ll}\text { - } & (\mathrm{RS})_{57} \\ \text { - } & (\mathrm{RS})_{60: 63} \\ \text { - } & (\mathrm{RB})_{37: 51}\end{array}\)
If this instruction is executed in 32-bit mode, \((R B)_{0: 31}\) must be zeros (i.e., the ESID must be in the range \(0: 15\) ).

This instruction must not be used to load a segment descriptor that is in the Segment Table when LPCR \(_{\mathrm{U}}\) PRT=1, and cannot be used to invalidate the translation contained in an SLB entry.

This instruction is privileged.

\section*{Special Registers Altered:}

None

\section*{Programming Note}

The reason slbmte must not be used to load segment descriptors that are in the Segment Table is that there could be a race condition with hardware loading the same segment descriptor, resulting in duplicate SLB entries. Software must not allow duplicate SLB entries to be created; see Section 5.7.8.2, "SLB Search".

The reason slbmte cannot be used to invalidate an SLB entry is that it does not necessarily affect implementation-specific address translation lookaside information. slbie (or slbia) must be used for this purpose.

SLB Move From Entry VSID
X-form
slbmfev RT,RB
\begin{tabular}{|c|c|c|c|c|c|c|c|c|}
\hline 31 & RT & III & L & RB & & 851 & \(/\) \\
\hline 0 & & 6 & & 11 & 15 & 16 & 21 & \\
31 \\
\hline
\end{tabular}

This instruction is used to read software-loaded SLB entries. When LPCR \({ }_{\text {UPRT }}=0\), the entry is specified by bits 52:63 of register RB. When LPCR first four entries can be read, so bits 52:61 of register RB are ignored. If the specified entry is valid ( \(\mathrm{V}=1\) ), the contents of the B, VSID, \(\mathrm{K}_{\mathrm{s}}, \mathrm{K}_{\mathrm{p}}, \mathrm{N}, \mathrm{L}, \mathrm{C}\), and LP fields of the entry are placed into register RT. The contents of these registers are interpreted as shown in Figure 48.
RT
\begin{tabular}{|l|ll|l|l|l|ll|}
\hline B & & VSID & \(\mathrm{K}_{\mathrm{s}} \mathrm{K}_{\mathrm{p}}\) NLC & 0 & LP & \multicolumn{2}{|c|}{ Os } \\
\hline 0 & 2 & & 52 & 57 & 58 & 60 & 63 \\
\hline
\end{tabular}

RB

\begin{tabular}{ll}
\(\mathrm{RT}_{0: 1}\) & B \\
\(\mathrm{RT}_{2: 51}\) & VSID \\
\(\mathrm{RT}_{52}\) & \(\mathrm{~K}_{\mathrm{s}}\) \\
\(\mathrm{RT}_{53}\) & \(\mathrm{~K}_{\mathrm{p}}\) \\
\(\mathrm{RT}_{54}\) & N \\
\(\mathrm{RT}_{55}\) & L \\
\(\mathrm{RT}_{56}\) & C \\
\(\mathrm{RT}_{57}\) & set to 0b0 \\
\(\mathrm{RT}_{58: 59}\) & LP \\
\(\mathrm{RT}_{60: 63}\) & set to \(0 b 0000\)
\end{tabular}
\(\mathrm{RB}_{0: 51} \quad\) must be \(0 \times 0 \_0000 \_0000 \_0000\)
\(\mathrm{RB}_{52: 63}\) index, which selects the SLB entry
Figure 48. GPR contents for slbmfev
On implementations that support a virtual address size of only n bits, \(\mathrm{n}<78, \mathrm{RT}_{2: 79-\mathrm{n}}\) are set to zeros.
If the SLB entry specified by bits 52:63 of register RB is invalid ( \(\mathrm{V}=0\) ), the contents of register RT are set to 0 .

High-order bits of \((\mathrm{RB})_{52: 63}\) that correspond to SLB entries beyond the size of the SLB provided by the implementation must be zeros.

The hardware ignores the contents of \(\mathrm{RB}_{0: 51}\).
This instruction is privileged.
The use of the \(L\) field is implementation specific.

\section*{Special Registers Altered:}

None

SLB Move From Entry ESID
X-form
slbmfee RT,RB
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline 31 & RT & I/I & L & RB & & 915 & \(/\) \\
\hline 0 & & & 11 & 15 & 16 & 21 & \\
31 \\
\hline
\end{tabular}

This instruction is used to read software-loaded SLB entries. When LPCR \({ }_{\text {UPRT }}=0\), the entry is specified by bits \(52: 63\) of register RB. When LPCR first four entries can be read, so bits 52:61 of register RB are ignored. If the specified entry is valid ( \(\mathrm{V}=1\) ), the contents of the ESID and V fields of the entry are placed into register RT. If \(L P C R_{U P R T}=1\), the value of the BO field of the entry is also placed into register RT. The contents of these registers are interpreted as shown in Figure 49.

RT
\begin{tabular}{|l|l|l|ll|}
\hline ESID & V & BO & Os \\
\hline 0 & 36 年 37 & 38 & 63
\end{tabular}

RB

\begin{tabular}{ll}
\(\mathrm{RT}_{0: 35}\) & ESID \\
\(\mathrm{RT}_{36}\) & V \\
\(\mathrm{RT}_{37}\) & BO, entry is bolted \\
\(\mathrm{RT}_{38: 63}\) & set to 0b000 II 0x00_0000 \\
\(\mathrm{RB}_{0: 51}\) & must be 0x0_0000_0000_0000 \\
\(\mathrm{RB}_{52: 63}\) & index, which selects the SLB entry
\end{tabular}

Figure 49. GPR contents for slbmfee
If the SLB entry specified by bits 52:63 of register RB is invalid ( \(\mathrm{V}=0\) ), the contents of register RT are set to 0 .

High-order bits of \((\mathrm{RB})_{52: 63}\) that correspond to SLB entries beyond the size of the SLB provided by the implementation must be zeros.

The hardware ignores the contents of \(\mathrm{RB}_{0: 51}\).
This instruction is privileged.
The use of the \(L\) field is implementation specific.

\section*{Special Registers Altered:}

None

\section*{SLB Find Entry ESID}

X-form
slbfee. RT,RB
\begin{tabular}{|c|c|c|c|c|c|c|}
\hline 31 & RT & & I/I & RB & & 979 \\
\hline 0 & & & 6 & 11 & 16 & 21 \\
\hline
\end{tabular}

The SLB is searched for an entry that matches the effective address specified by register RB. When LPCR \(_{\text {UPRT }}=1\), this instruction is nonfunctional. The search is performed as if it were being performed for purposes of address translation. That is, in order for a given entry to satisfy the search, the entry must be valid \((V=1)\), and \((R B)_{0: 63-s}\) must equal SLBE \(\left[E S I D_{0: 63-s}\right]\) (where \(2^{s}\) is the segment size selected by the \(B\) field in the entry). If exactly one matching entry is found, the contents of the B, VSID, \(K_{s}, K_{p}, N, L, C\), and LP fields of the entry are placed into register RT. If no matching entry is found, register RT is set to 0 . If more than one matching entry is found, either one of the matching entries is used, as if it were the only matching entry, or a Machine Check occurs. If a Machine Check occurs, register RT, and CR Field 0 are set to undefined values, and the description below of how this register and this field is set does not apply.

The contents of registers RT and RB are interpreted as shown in Figure 50.

RT
\begin{tabular}{|l|ll|l|l|l|l|l|}
\hline B & VSID & \(\mathrm{K}_{\mathrm{s}} \mathrm{K}_{\mathrm{p}}\) NLC & O & LP & & Os & \\
\hline 0 & 2 & 52 & 57 & 58 & 60 & & 63 \\
\hline
\end{tabular}

RB
\begin{tabular}{|l|l|ll|}
\hline ESID & 0000 & Os \\
\hline 0 & 36 & 40 & 63 \\
\hline
\end{tabular}
\begin{tabular}{ll}
\(\mathrm{RT}_{0: 1}\) & B \\
\(\mathrm{RT}_{2: 51}\) & VSID \\
\(\mathrm{RT}_{52}\) & \(\mathrm{~K}_{\mathrm{s}}\) \\
\(\mathrm{RT}_{53}\) & \(\mathrm{~K}_{\mathrm{p}}\) \\
\(\mathrm{RT}_{54}\) & N \\
\(\mathrm{RT}_{55}\) & L \\
\(\mathrm{RT}_{56}\) & C \\
\(\mathrm{RT}_{57}\) & set to 0b0 \\
\(\mathrm{RT}_{58: 59}\) & LP \\
\(\mathrm{RT}_{60: 63}\) & set to 0b0000 \\
\(\mathrm{RB}_{0: 35}\) & ESID \\
\(\mathrm{RB}_{36: 39}\) & must be \(0 b 0000\) \\
\(\mathrm{RB}_{40: 63}\) & must be \(0 \times 000000\)
\end{tabular}

Figure 50. GPR contents for slbfee.
If \(s>28, \mathrm{RT}_{80-\mathrm{s}: 51}\) are set to zeros. On implementations that support a virtual address size of only n bits, n \(<78, \mathrm{RT}_{2: 79-n}\) are set to zeros.
CR Field 0 is set as follows. \(j\) is a 1 -bit value that is equal to Ob1 if a matching entry was found. Otherwise, \(j\) I
\[
\mathrm{CRO}_{\text {LT GT EQ }} \mathrm{so}=0 \mathrm{~b} 00\|\mathrm{l}\| \mathrm{XER}_{\mathrm{SO}}
\]

The hardware ignores the contents of \(\mathrm{RB}_{36: 38} 40: 63\).
If this instruction is executed in 32-bit mode, \((\mathrm{RB})_{0: 31}\) must be zeros (i.e., the ESID must be in the range 0-15).

This instruction is privileged.
```

Special Registers Altered:
CRO

```

\section*{SLB Synchronize \\ X-form}
slbsync


The slbsync instruction provides an ordering function for the effects of all slbieg instructions executed by the thread executing the slbsync instruction, with respect to the memory barrier created by a subsequent ptesync instruction executed by the same thread. Executing a slbsync instruction ensures that all of the following will occur.
- All SLB invalidations caused by slbieg instructions preceding the slbsync instruction will have completed on any other thread before any data accesses caused by instructions following the ptesync instruction are performed with respect to that thread.
- All storage accesses by other threads for which the address was translated using the translations being invalidated will have been performed with respect to the thread executing the ptesync instruction, to the extent required by the associated Memory Coherence Required attributes, before the ptesync instruction's memory barrier is created.

The operation performed by this instruction is ordered by the eieio (or sync or ptesync) instruction with respect to preceding slbieg instructions executed by the thread executing the slbsync instruction. The operations caused by slbieg and slbsync are ordered by eieio as a fifth set of operations, which is independent of the other four sets that eieio orders.

The slbsync instruction may complete before operations caused by slbieg instructions preceding the slbsync instruction have been performed.
This instruction is privileged except when \(\mathrm{LPCR}_{\text {GTSE }}=0\), making it hypervisor privileged.
See Section 5.10 for a description of other requirements associated with the use of this instruction.

\section*{Special Registers Altered:}

None

Version 3.0

I Programming Note
slbsync should not be used to synchronize the completion of slbie.

\subsection*{5.9.3.3 TLB Management Instructions}

In addition to managing the TLB, tlbie and tlbiel are also used to manage the Page Walk Cache, In-Memory Table caching, and implementation-specific lookaside information that depends on the values of the PTEs. The parameters described below specify the type of translations to invalidate and the scope of the invalidation to be performed.
Radix Invalidation Control (RIC) specifies whether to invalidate the TLB, the Page Walk Cache, or both together with partition and Process Table caching. The RIC values and functions are as follows.

0 Just invalidate TLB.
1 Invalidate just Page Walk Cache.
2 Invalidate TLB, Page Walk Cache, and any caching of Partition and Process Table Entries.

3 Invalidate a series of consecutive translations (just in the TLB).

Process Scoped (PRS) specifies whether the translation(s) to be invalidated are partition scoped or process scoped including, for RIC=2, whether process or Partition Table caching is being invalidated.
\(0 \quad\) Invalidate partition-scoped translation(s).
1 Invalidate process-scoped translations.
Radix ( R ) specifies whether the translations to be invalidated are Radix Tree translations or HPT translations.
\(0 \quad\) Invalidate HPT translation(s).
1 Invalidate Radix Tree translations.
Invalidation Selector (IS) (found in RB) specifies the scope of the context to be invalidated.

0 Invalidate just the target VA.
1 Invalidate matching PID.
2 Invalidate matching LPID.
3 If \(M S R_{H V}=1\), invalidate all entries, otherwise invalidate matching LPID.

The IS \(\neq 0\) RIC=2 variants of \(\boldsymbol{t} / \boldsymbol{b i e}\) and \(\boldsymbol{t} / \boldsymbol{b i e l}\) perform the same TLB invalidations as the corresponding \(\mathrm{RIC}=0\) variants, but in addition invalidate Page Walk Cache Entries and partition or Process Table caching associated with the specified LPID or LPID/PID. When RIC=1 and \(I S \neq 0\), the Page Walk Cache Entries for the specified LPID or LPID/PID are invalidated while leaving the corresponding TLB entries intact. The ability to target an individual Page Walk Cache Entry or the set of entries associated with a given Page Table Entry (i.e. IS=0 for RIC=1 or RIC=2) is not supported by the Power ISA. When RIC=3 and IS=0, t/bie invalidates a series of consecutive translations or translations associated with an aligned region of address space for HPT
translation. The \(\mathrm{IS} \neq 0\) tlbiel variants operate on a specified congruence class, requiring a software loop where tlbie operates on the entire TLB. For IS=0 invalidations of Radix Tree translations, the use of tlbie[I] is limited to translations for quadrant 0 .

When reassigning an LPID or PID, after updating the Partition and/or Process Table(s) software must use a tlbie instruction to remove lookaside information associated with the old parition or process.

To invalidate TLB entries, software must supply an effective page number for process-scoped Radix Tree translations, a guest real page number for parti-tion-scoped Radix Tree translations, and an abbreviated virtual page number for HPT translations. The RTL, RB illustration, and verbal description for \(R=1\) require the reader to make the appropriate mental substitution for partition-scoped invalidation. Note also that where page size is specified to be a function of \(L\) and AP, it may also be a function of \(L\) and LP. The architecture allows for three independent sets of page sizes, one for \(\mathrm{R}=1\), one for RIC=3 (requires \(\mathrm{R}=0\) ), and one for all other cases. An implementation may choose to have a single set of encodings work consistenty between any two or all three states.

\section*{Programming Note}

Changes to the Page Table in the presence of active transactions may compromise transactional semantics if a page accessed by a translation is remapped within the lifetime of the transaction. Through the use of a tlbie instruction to the unmapped page, an operating system or hypervisor can ensure that any transaction that has touched the affected page is terminated.
Changes to local translation lookaside buffers, through the tlbiel instruction, have no effect on transactions. Consequently, if these instructions are used to invalidate TLB entries after the unmapping of a page, it is the responsibility of the OS or hypervisor to ensure that any transaction that may have touched the modified page is terminated, using a tabort. or treclaim instruction.

\section*{TLB Invalidate Entry}

\section*{X-form}
\begin{tabular}{|c|c|c|c|c|c|c|c|}
\hline \multicolumn{8}{|l|}{tlbie RB,RS,RIC,PRS,R} \\
\hline \(0^{31}\) & \({ }_{6} \mathrm{RS}\) & \begin{tabular}{|l|l|}
\hline 11 & RIC \\
12
\end{tabular} & PRS & \begin{tabular}{|c|c} 
R & RB \\
15 & 16 \\
\hline
\end{tabular} & 21 & 306 & 1
31 \\
\hline
\end{tabular}
```

IS }\leftarrow(\textrm{RB}\mp@subsup{)}{52:53}{
else search_LPID=LPIDR LPID
switch(IS)
case (0b00):
If RIC=0
if R=0 then
L}\leftarrow(RB)6
if L = 0
then

```
if \(\mathrm{MSR}_{\mathrm{HV}}=1\) then search_LPID=RS \(32: 63\)
                    base_pg_size \(=4 \mathrm{~K}\)
                    actual_pg_size =
                        page size specified in (RB) 56:58
                    \(i=51\)
            else
                    base_pg_size =
                            base page size specified in (RB) \(44: 51\)
                    actual_pg_size =
                    actual page size specified in (RB) \(44: 51\)
                    \(\mathrm{b} \leftarrow\) log_base_2 (base_pg_size)
                    \(p \leftarrow \log _{\text {_base_2 }}\) (actual_pg_size)
                    \(i=\max (\min (43,63-b), 63-\mathrm{p})\)
        sg_size \(\leftarrow\) segment size specified in \((R B)_{54: 55}\)
        for each thread
            for each TLB entry
                    if (entry_ \(\left.\mathrm{VA}_{14: i+14}=(\mathrm{RB})_{0: i}\right)\) \&
                    (entry_sg_size = sg_size) \&
                        (entry_base_pg_size = base_pg_size) \&
                            (entry_actual_pg_size =
                        actual_pg_size) \&
                        (entry_LPID = search_LPID) \&
                        (entry_process_scoped = 0) \&
                        (entry_radix = 0)
                then
                if \(((L=0) \mid(b \geq 20))\) then
                        TLB entry \(\leftarrow\) invalid
                else
                if (entry_ \(\mathrm{VA}_{58: 77-\mathrm{b}}=(\mathrm{RB})_{56: 75-\mathrm{b}}\) ) then
                    TLB entry \(\leftarrow\) invalid
        else
        actual_pg_size =
                page size specified in (RB) 56:58
            \(\mathrm{p} \leftarrow\) log_base_2 (actual_pg_size)
            \(i=63-p\)
            for each thread
                for each TLB entry
                    if (entry_EA \(\left.A_{0: i}=(R B)_{0: i}\right)\) \&
                        (entry_actual_pg_size =
                        actual_pg_size) \&
                    (entry_LPID = search_LPID) \&
                    (entry_process_scope \(\bar{d}=\) PRS) \&
                    (entry_radix = R) \&
                    \(((\operatorname{PRS}=0) \mid\)
                        \(\left(\right.\) entry_PID \(\left.\left.=(R S)_{0: 31}\right)\right)\)
                    then
                        TLB entry \(\leftarrow\) invalid
    else if RIC=3 then
```

        if RB}\mp@subsup{\textrm{AP}}{\mathrm{ P }}{}\mathrm{ indicates cluster bomb then
        n = implementation-specific series size
            f(L||AP)
        base_pg_size =
        implementation-specific base page
            size f(L||AP)
    trunc= log}2(n * base_pg_size)-1
    loop_RB }\leftarrow R
    loop_R\mp@subsup{B}{AP}{}}\leftarrow implementation-specific
                                    encoding for base_pg_size
    loop_RB}\mp@subsup{B}{\mathrm{ VPN }}{}\leftarrow loop_RB VPN[0:51-trunc] || trunc 0 
    do i=0 to n-1
        tlbie loop_RB,RS,0,0,0
        loop_RB VPN}\leftarrowloop_RB VPN + (base_pg_size/4096)
    else /* range bomb */
    range = implementation-specific range size
                f(L||AP)
    trunc= 増(range) -12
    loop_RB \leftarrow RB
    loop_RB VPN }\leftarrow loop_RB BPN[0:51-trunc] || trunc 0 
    for each TLB entry for each thread
    ```

```

            (entry_LPID = search_LPID) then
            TLB entry \leftarrow invalid
    case (0b01):
if RIC=0 | RIC=2 then
for each TLB entry for each thread
if (entry_LPID=search_LPID)
\&(entry_PID=RS0:31)
\&(entry_R=1)
\&(entry_PRS=1)
then TLB entry \leftarrow invalid
if RIC=1 | RIC=2 then
for each thread
invalidate process-scoped radix page walk
caching associated with process RS 0:31 in
partition search_LPID
if (RIC=2)\&(PRS=1) then
for each thread
invalidate Process Table caching associated
with process RS 0:31 in partition search_LPID
case (0b10):
if RIC=0 | RIC=2 then
if (PRS=0)\&(MSR
for each partition-scoped TLB entry for each
thread
if entry_LPID=search_LPID
then TLB entry \leftarrow invalid
if PRS=1 then
for each process-scoped TLB entry for each
thread
if entry_LPID=search_LPID
then TLB entry \leftarrow invalid
if RIC=1 | RIC=2 then
for each thread
if (PRS=0)\& (MSR
for each thread invalidate partition-
scoped page walk caching associated with
partition search_LPID
if PRS=1 then
for each thread invalidate process-scoped
page walk caching associated with
partition search_LPID
if RIC=2 then
if (PRS=0)\&(MSR

```
```

            for each thread invalidate Partition Table
            caching associated with partition
            search_LPID
    if PRS=1 then
            for each thread invalidate Process Table
            caching associated with partition
            search_LPID
    case (0.b11):
if RIC=0 | RIC=2 then
if MSR
for all threads
if PRS=0 then
all partition-scoped TLB entries
\leftarrowinvalid
else
all process-scoped TLB entries \leftarrowinvalid
if (MSR
for each process-scoped TLB entry for each
thread
if TLBE
then TLB entry \leftarrow invalid
if RIC=1 | RIC=2 then
if MSR HV then
if PRS=0 then
for all threads
invalidate all partition-scoped
page walk caching
else
for all threads
invalidate all process-scoped
page walk caching
if (MSR
for each thread invalidate process-scoped
page walk caching associated with
partition search_LPID
if RIC=2 then
if MSR
if PRS=0 then
for each thread
invalidate all Partition Table caching
else
for each thread
invalidate all Process Table caching
if (MSR
for each thread invalidate Process Table
caching associated with partition
search_LPID

```

The operation performed by this instruction is based on the contents of registers RS and RB. The contents of these registers are shown below, where IS is \((\mathrm{RB})_{52: 53}\) and \(L\) is \((R B)_{63}\).
RS:

| RB for R=1 and IS=0b00::


RB for \(R=0, I S=0 b 00\), and \(L=0::\)
\begin{tabular}{|c|c|c|c|c|c|}
\hline AVA & IS & B & AP & Os & L \\
\hline 0 & 52 & 54 & 56 & 59 & 63 \\
\hline
\end{tabular}

RB for \(R=0\), \(I S=0 b 00\), and \(L=1\) :
\begin{tabular}{|l|l|l|l|l|l|}
\hline AVA & LP & IS & B & AVAL & L \\
\hline 0 & 44 & 52 & 54 & 56 & 63 \\
\hline
\end{tabular}

RB for IS=0b01, Ob10, or Ob11:
\begin{tabular}{|l|l|ll|}
\hline & Os & IS & \multicolumn{2}{|c|}{ Os } \\
\hline 0 & 52 & 54 & \\
\hline
\end{tabular}

If this instruction is executed in hypervisor state, \(\mathrm{RS}_{32: 63}\) contains the partion ID (LPID) of the partition for which one or more translations are being invalidated. Otherwise, the value in LPIDR is used. The supported \((\mathrm{RS})_{32: 63}\) values are the same as the LPID values supported in LPIDR. \(\mathrm{RS}_{0: 31}\) contains a PID value. The supported values of \(\mathrm{RS}_{0: 31}\) are the same as the PID values supported in PIDR.

The following forms are invalid.
- \(\mathrm{PRS}=1, \mathrm{R}=0\), and \(\mathrm{RIC} \neq 2\) (The only pro-cess-scoped HPT caching is of the Process Table.)
- RIC=1 and \(\mathrm{R}=0\) (There is no Page Walk Cache for HPT translation.)
- RIC=3 and \(\mathrm{R}=1\) (Cluster bombs are only supported for HPT translation.)

The following forms are treated as if the instruction form were invalid.
- RIC=1 and IS=0 (The architecture does not support shootdown of individual translations in the Page Walk Cache.)
- RIC=2 and IS=0 (RIC is for comprehensive invalidation that is not supported at the level of an individual page.)
- RIC=3 and \(I S \neq 0\) (Cluster bombs are only supported for individual pages.)
- PRS=0 and IS=1 (Partition-scoped translations are not associated with processes.)
- \(\mathrm{R}=0\) and \(\mathrm{IS}=1\) (HPT translations are not associated with processes.)
The results of an attempt to invalidate a translation outside of quadrant 0 for Radix Tree translation ( \(R=1\), RIC \(=0\), \(P R S=1, I S=0\), and \(E A_{0: 1} \neq 0 b 00\) ) are boundedly undefined.

\section*{IS field in RB contains Ob00}

If RIC=0, this is a search for a single TLB entry. The following relationships must be true and tests and actions are performed to search for an HPT translation.

If the base page size specified by the PTE that was used to create the TLB entry to be invalidated is 4 \(K B\), the \(L\) field in register RB must contain 0 .

If the \(L\) field in RB contains 0 , the base page size is 4 KB and \(\mathrm{RB}_{56: 58}\) (AP - Actual Page size field)
must be set to the SLBE \(_{\text {LIILP }}\) encoding for the page size corresponding to the actual page size specified by the PTE that was used to create the TLB entry to be invalidated. Thus, \(b\) is equal to 12 and \(p\) is equal to \(\log _{2}\) (actual page size specified by (RB) \({ }_{56: 58}\) ). The Abbreviated Virtual Address (AVA) field in register RB must contain bits 14:65 of the virtual address translated by the TLB entry to be invalidated. Variable i is equal to 51.
If the \(L\) field in RB contains 1, the following rules apply.
■ The base page size and actual page size are specified in the LP field in register RB, where the relationship between (RB) 44:51 (LP - Large Page size selector field) and the base page size and actual page size is the same as the relationship between \(P T E_{\text {LP }}\) and the base page size and actual page size, except for the " \(r\) " bits (see Section 5.7.9.1 on page 998 and Figure 31 on page 999). Thus, b is equal to \(\log _{2}\) (base page size specified by \((R B)_{44: 51}\) ) and \(p\) is equal to \(\log _{2}\) (actual page size specified by (RB) \({ }_{44: 51}\) ). Specifically, (RB) \({ }_{44+c: 51}\) must be equal to the contents of bits \(\mathrm{c}: 7\) of the LP field of the PTE that was used to create the TLB entry to be invalidated, where \(c\) is the maximum of 0 and ( \(20-\mathrm{p}\) ).
- Variable i is the larger of ( \(63-\mathrm{p}\) ) and the value that is the smaller of 43 and (63-b). (RB) \(0: i\) must contain bits \(14:(i+14)\) of the virtual address translated by the TLB to be invalidated. If \(b>20, \mathrm{RB}_{64-\mathrm{b}: 43}\) may contain any value and are ignored by the hardware.
■ If \(b<20\), (RB) \({ }_{56: 75-b}\) must contain bits 58:77-b of the virtual address translated by the TLB to be invalidated, and other bits in (RB) \(56: 62\) may contain any value and are ignored by the hardware.
■ If \(b \geq 20,(R B)_{56: 62}\) (AVAL - Abbreviated Virtual Address, Lower) may contain any value and are ignored by the hardware.

Let the segment size be equal to the segment size specified in (RB) \({ }_{54: 55}\) (B field). The contents of \(\mathrm{RB}_{54: 55}\) must be the same as the contents of the \(B\) field of the PTE that was used to create the TLB entry to be invalidated.
\(\mathrm{RB}_{52: 53}\) and \(\mathrm{RB}_{59: 62}\left(\right.\) when \((\mathrm{RB})_{63}=0\) ) must contain zeros and are ignored by the hardware.
All TLB entries on all threads that have all of the following properties are made invalid.
- The entry translates a virtual address for which all the following are true.
- \(\mathrm{VA}_{14: 14+\mathrm{i}}\) is equal to \((R B)_{0: i}\).
- \(L=0\) or \(b \geq 20\) or, if \(L=1\) and \(b<20\), \(\mathrm{VA}_{58: 77-\mathrm{b}}\) is equal to (RB) \(56: 75\)-b
\(\square\) The segment size of the entry is the same as the segment size specified in (RB) 54:55 .
- Either of the following is true:
- The \(L\) field in RB is 0 , the base page size of the entry is 4 KB , and the actual page size of the entry matches the actual page size specified in (RB) 56:58 \(^{5}\).
- The \(L\) field in RB is 1 , the base page size of the entry matches the base page size specified in (RB) 44:51 , and the actual page size of the entry matches the actual page size specified in (RB) \({ }_{44: 51}\).
- TLBE \(_{\text {LPID }}=\) search_LPID.

Additional TLB entries may also be made invalid if those TLB entries contain an LPID that matches search_LPID.

The following relationships must be true and tests and actions are performed to search for a Radix Tree translation. For a partition-scoped invalidation, references to the effective address are understood to refer to the guest real address.

The page size is encoded in \(\mathrm{RB}_{56: 58}\) (AP - Actual Page size field). Thus \(p\) is equal to \(\log _{2}\) ( page size specified by \(\mathrm{RB}_{56: 58}\) ). The Effective Page Number (EPN) field in register RB must contain the bits 0:i of the effective address translated by the TLB entry to be invalidated. Variable \(i\) is equal to \(63-\mathrm{p}\).

The fields shown as zeros must be set to zero and are ignored by the hardware.
All TLB entries on all threads that have all of the following properties are made invalid.
- The entry translates an effective address for which \(E A_{0: i}\) is equal to \((R B)_{0: i}\).
- The page size of the entry matches the page size specified in (RB) 56:58 .
■ The entry has the appropriate scope (partition or process).
■ The process ID specified in RS matches the process ID in the TLB entry if not invalidating a partition-scoped translation.
- TLBE \({ }_{\text {LPID }}\) matches the partiion ID of the partition for which the translation is to be invalidated.
Additional TLB entries may also be made invalid if those TLB entries contain an LPID that matches the partition ID of the partition for which the translation is to be invalidated.
If RIC=3, then an implementation-specific encoding of AP indicates the number and (base) size of a series of sequential virtual pages or an address range for which the translations will be invalidated. The range or pages occupy an aligned region of virtual storage. The address in RB is masked to get the base address of the region. For the cluster bomb version, tlbies are then performed for the virtual pages within the region. For the range bomb version, the entire TLB is searched and all entries that map a portion of the address range for the target partition are marked invalid.

\section*{IS field in RB is non-zero}

If RIC=0 or RIC=2, all partition-scoped TLB entries when PRS=0 and \(M S R_{H V}=1\) or all process-scoped TLB entries when PRS \(=1\) on all threads for which any of the following conditions are met for the entry are made invalid.
- The IS field in RB contains \(0 b 10\) or \(\mathrm{MSR}_{\mathrm{HV}}=0\) and the IS field contains 0b11, and TLBE LPID matches the partition ID of the partition for which the translation is to be invalidated.
■ The IS field in RB contains Ob01, TLBE matches the partition ID of the partition for which the translation is to be invalidated, and TLBE \(_{\text {PID }}=\mathrm{RS}_{0: 31}\).
- The IS field in RB contains \(0 b 11\) and \(M S R_{H V}=1\).

If \(\mathrm{RIC}=1\) or \(\mathrm{RIC}=2\), if the following conditions are met, the respective partition-scoped contents when \(\mathrm{PRS}=0\) and \(\mathrm{MSR}_{\mathrm{HV}}=1\) or process-scoped contents when \(\mathrm{PRS}=1\) of the page walk cache are invalidated.
- If the IS field in RB contains \(0 b 10\) or if IS contains \(0 b 11\) and \(\mathrm{MSR}_{\mathrm{HV}}=0\), for all threads, all prop-erly-scoped page walk caching associated with the partition for which the translation is to be invalidated is invalidated.
■ If the IS field in RB contains 0b11 and \(\mathrm{MSR}_{\mathrm{HV}}=1\), the entire properly-scoped page walk caching for each thread is invalidated.
- If the IS field in RB contains 0b01 (and PRS=1 and \(R=1\) ), for all threads, all properly-scoped page walk caching associated with process \(\mathrm{RS}_{0: 31}\) in the partition for which the translation is to be invalidated is invalidated.
If RIC=2, if the following conditions are met, the respective partition and Process Table caching are invalidated for all threads.
- If the IS field in RB contains \(0 b 01\) and PRS=1, for all threads, caching of Process Table Entries for process \(\mathrm{RS}_{0: 31}\) in the partition for which the translation is to be invalidated is invalidated.
- If the IS field in RB contains 0b10, \(M_{S R}=1\), and PRS=0, for all threads, caching of Partition Tables for the partition for which the translation is to be invalidated is invalidated.
- If the IS field in RB contains \(0 b 10\) and PRS=1, for all threads, caching of Process Tables for the partition for which the translation is to be invalidated is invalidated.
■ if the IS field in RB contains 0b11, MSR \(\mathrm{MV}_{\mathrm{HV}}=1\), and PRS=0, for all threads, all Partition Table caching is invalidated.
- if the IS field in RB contains 0b11, \(M_{S R} R_{H V}=1\), and PRS=1, for all theads, all Process Table caching is invalidated.
- If the IS field in RB contains 0b11, \(\mathrm{MSR}_{H V}=0\), and PRS=1, for all threads, caching of Process Tables for the partition for which the translation is to be invalidated is invalidated.

When \(\mathrm{i}>40, \mathrm{RB}_{40: i-1}\) may contain any value and are ignored by the hardware.

\section*{For all IS values}

For all threads, any implementation specific lookaside information that is based on any TLB entry that would be invalidated by this instruction will also be invalidated.
\(\mathrm{MSR}_{\mathrm{SF}}\) must be 1 when this instruction is executed; otherwise the results are undefined.

If the value specified in \(\mathrm{RS}_{0: 31}, \mathrm{RS}_{32: 63}, \mathrm{RB}_{54: 55}\) when \(R=0, \mathrm{RB}_{56: 58}\) when \(\mathrm{RB}_{63}=0\), or \(\mathrm{RB}_{44: 51}\) when \(\mathrm{RB}_{63}=1\) is not supported by the implementation, the instruction is treated as if the instruction form were invalid.

The operation performed by this instruction is ordered by the eieio (or sync or ptesync) instruction with respect to a subsequent tlbsync instruction executed by the thread executing the tlbie instruction. The operations caused by tlbie and tlbsync are ordered by eieio as a fourth set of operations, which is independent of the other three sets that eieio orders.

This instruction is privileged except when \(\operatorname{LPCR}_{\text {GTSE }}=0\) or when \(\mathrm{PRS}=0\) and HR\|GR \(\neq 0 \mathrm{~b} 00\), making it hypervisor privileged.
See Section 5.10, "Translation Table Update Synchronization Requirements" for a description of other requirements associated with the use of this instruction.
Special Registers Altered:
None

\section*{Extended Mnemonics:}

Extended mnemonic for tlbie::

\section*{Extended:}
tlbie RB,RS

Equivalent to:
tlbie RB,RS,0,0,0

\section*{Special Registers Altered: \\ None}

\section*{Programming Note}

For tlbie[ \(I]\) instructions in which (RB) \()_{63}=0\), the AP value in RB is provided to make it easier for the hardware to locate address translations, in lookaside buffers, corresponding to the address translation being invalidated.
For t/bie[ \(/\) instructions the AP specification is not binary compatible with versions of the architecture that precede Version 2.06. As an example, for an actual page size of \(64 \mathrm{~KB} A P=0 b 101\), whereas software written for an implementation that complies with a version of the architecture that precedes V. 2.06 would have AP=100 since AP was a 1 bit value followed by 0 s in \(\mathrm{RB}_{57: 58}\). If binary compatibility is important, for a 64 KB page software can use AP=0b101 on these earlier implementations since these implementations were required to ignore \(\mathrm{RB}_{57: 58}\).

\section*{Programming Note}

For tlbie[ \(I\) ] instructions the AVA and AVAL fields in RB contain different VA bits from those in PTE \(_{\text {AVA }}\).
TLB Invalidate Entry Local
                            X-form
tlbiel RB,RS,RIC,PRS,R
\begin{tabular}{|c|c|c|l|l|l|l|l|r|r|}
\hline 31 & RS & RIC & PRS & R & RB & & 274 & \(/\) \\
0 & & 6 & 11 & 12 & 14 & 15 & 16 & 21 & \\
31 \\
\hline
\end{tabular}
IS \(\leftarrow(\mathrm{RB})_{52: 53}\)
search_LPID=LPIDR \({ }_{\text {LPID }}\)
switch (IS)
    case (0b00):
        If \(\mathrm{RIC}=0\)
        If \(\mathrm{R}=0\)
                \(\mathrm{L} \leftarrow(\mathrm{RB})_{63}\)
        if \(L=0\) then
            base_pg_size \(=4 \mathrm{~K}\)
            actual_pg_size =
                    page size specified in (RB) 56:58
            i = 51
        else
            base_pg_size = base page size specified
                in ( RB ) \(44: 51\)
            actual_pg_size =
                actual page size specified in (RB) 44:51
            \(\mathrm{b} \leftarrow\) log_base_2 (base_pg_size)
            \(p \leftarrow \log\) _base_2 (actual_pg_size)
            \(i=\max (\min (43,63-b), 63-\mathrm{p})\)
        sg_size \(\leftarrow\) segment size specified in \((R B)_{54: 55}\)
        for each TLB entry
            if (entry_ \(\left.\mathrm{VA}_{14: i+14}=(\mathrm{RB})_{0: i}\right)\) \&
                (entry_sg_size = segment_size) \&
                (entry_base_pg_size = base_pg_size) \&
                (entry_actual_pg_size =actual_pg_size) \&
| else
```

```
```

        (TLBE 
    ```
```

```
        (TLBE 
```

```
```

        (TLBE 
        (entry_process_scoped=0) &
        (entry_process_scoped=0) &
        (entry_process_scoped=0) &
        (entry_radix=0))
        (entry_radix=0))
        (entry_radix=0))
        then
        then
        then
        if ((L = 0)|(b\geq20)) then
        if ((L = 0)|(b\geq20)) then
        if ((L = 0)|(b\geq20)) then
                        TLB entry \leftarrow invalid
                        TLB entry \leftarrow invalid
                        TLB entry \leftarrow invalid
        else
        else
        else
        if (entry_VA58:77-b}=(RB) 56:75-b) then
        if (entry_VA58:77-b}=(RB) 56:75-b) then
        if (entry_VA58:77-b}=(RB) 56:75-b) then
                TLB entry \leftarrow invalid
                TLB entry \leftarrow invalid
                TLB entry \leftarrow invalid
        pg_size = page size specified in (RB) 56:58
        pg_size = page size specified in (RB) 56:58
        pg_size = page size specified in (RB) 56:58
        p}\leftarrow log_base_2(pg_size
        p}\leftarrow log_base_2(pg_size
        p}\leftarrow log_base_2(pg_size
        i = 63-p
        i = 63-p
        i = 63-p
        for each TLB entry
        for each TLB entry
        for each TLB entry
            if (entry_EA 0:i = (RB) 0:i) &
            if (entry_EA 0:i = (RB) 0:i) &
            if (entry_EA 0:i = (RB) 0:i) &
                (entry_pg_size = pg_size) &
                (entry_pg_size = pg_size) &
                (entry_pg_size = pg_size) &
            (entry_LPID = search_LPID) &
            (entry_LPID = search_LPID) &
            (entry_LPID = search_LPID) &
                (entry_process_scoped = PRS) &
                (entry_process_scoped = PRS) &
                (entry_process_scoped = PRS) &
                (entry_radix = R) &
                (entry_radix = R) &
                (entry_radix = R) &
                ((PRS = 0) )
                ((PRS = 0) )
                ((PRS = 0) )
                (entry_PID = (RS) 0:31))
                (entry_PID = (RS) 0:31))
                (entry_PID = (RS) 0:31))
            then
            then
            then
            TLB entry \leftarrow invalid
            TLB entry \leftarrow invalid
            TLB entry \leftarrow invalid
    case (0b01):
case (0b01):
case (0b01):
if RIC=0 | RIC=2 then
if RIC=0 | RIC=2 then
if RIC=0 | RIC=2 then
i}\leftarrow implementation-dependent number, 40\leqi\leq51
i}\leftarrow implementation-dependent number, 40\leqi\leq51
i}\leftarrow implementation-dependent number, 40\leqi\leq51
for each TLB entry in set (RB) i:51
for each TLB entry in set (RB) i:51
for each TLB entry in set (RB) i:51
if (entry_LPID=search_LPID)
if (entry_LPID=search_LPID)
if (entry_LPID=search_LPID)
\&(entry_PID=RSS:31)
\&(entry_PID=RSS:31)
\&(entry_PID=RSS:31)
\&(entry_R=1)
\&(entry_R=1)
\&(entry_R=1)
\&(entry_PRS=1)
\&(entry_PRS=1)
\&(entry_PRS=1)
then TLB entry }\leftarrow invali
then TLB entry }\leftarrow invali
then TLB entry }\leftarrow invali
if RIC=1 | RIC=2 then
if RIC=1 | RIC=2 then
if RIC=1 | RIC=2 then
invalidate process-scoped radix page walk
invalidate process-scoped radix page walk
invalidate process-scoped radix page walk
caching associated with process RS 0:31 in
caching associated with process RS 0:31 in
caching associated with process RS 0:31 in
partition search_LPID
partition search_LPID
partition search_LPID
if (RIC=2)\&(PRS=1) then
if (RIC=2)\&(PRS=1) then
if (RIC=2)\&(PRS=1) then
invalidate Process Table caching associated
invalidate Process Table caching associated
invalidate Process Table caching associated
with process RS 0:31 in partition search_LPID
with process RS 0:31 in partition search_LPID
with process RS 0:31 in partition search_LPID
case (0b10):
case (0b10):
case (0b10):
if RIC=0 | RIC=2 then
if RIC=0 | RIC=2 then
if RIC=0 | RIC=2 then
i}~\mathrm{ implementation-dependent number, 40<i<51
i}~\mathrm{ implementation-dependent number, 40<i<51
i}~\mathrm{ implementation-dependent number, 40<i<51
if (PRS=0)\&(MSR
if (PRS=0)\&(MSR
if (PRS=0)\&(MSR
for each partition-scoped TLB entry in set
for each partition-scoped TLB entry in set
for each partition-scoped TLB entry in set
(RB) i:51
(RB) i:51
(RB) i:51
if entry_LPID=search_LPID
if entry_LPID=search_LPID
if entry_LPID=search_LPID
then TLB entry \leftarrow invalid
then TLB entry \leftarrow invalid
then TLB entry \leftarrow invalid
if PRS=1 then
if PRS=1 then
if PRS=1 then
for each process-scoped TLB entry in
for each process-scoped TLB entry in
for each process-scoped TLB entry in
set (RB) i:51
set (RB) i:51
set (RB) i:51
if entry_LPID=search_LPID
if entry_LPID=search_LPID
if entry_LPID=search_LPID
then TLB entry \leftarrow invalid
then TLB entry \leftarrow invalid
then TLB entry \leftarrow invalid
if RIC=1 | RIC=2 then
if RIC=1 | RIC=2 then
if RIC=1 | RIC=2 then
for each thread
for each thread
for each thread
if (PRS=0)\&(MSRHV=1) then
if (PRS=0)\&(MSRHV=1) then
if (PRS=0)\&(MSRHV=1) then
invalidate partition-scoped page walk
invalidate partition-scoped page walk
invalidate partition-scoped page walk
caching associated with partition
caching associated with partition
caching associated with partition
search_LPID
search_LPID
search_LPID
if PRS=1 then
if PRS=1 then
if PRS=1 then
invalidate process-scoped page walk
invalidate process-scoped page walk
invalidate process-scoped page walk
caching associated with partition
caching associated with partition
caching associated with partition
search_LPID
search_LPID
search_LPID
if RIC=2 then
if RIC=2 then
if RIC=2 then
if (PRS=0)\& (MSR
if (PRS=0)\& (MSR
if (PRS=0)\& (MSR
invalidate Partition Table caching
invalidate Partition Table caching
invalidate Partition Table caching
associated with partition search_LPID

```
```

            associated with partition search_LPID
    ```
```

            associated with partition search_LPID
    ```
```

```
        1f ((L - O) | (b \geq20)) then
```

```
        1f ((L - O) | (b \geq20)) then
```

```
        1f ((L - O) | (b \geq20)) then
```

```
        if PRS=1 then
            invalidate Process Table caching associated
            with partition search_LPID
case (0.b11):
| if RIC=0 | RIC=2 then
    i\leftarrowimplementation-dependent number, 40\leqi\leq51
    if MSR HV then
            if PRS=0 then
                all partition-scoped TLB entries in
                set (RB)i:51 \leftarrowinvalid
            else
                all process-scoped TLB entries in
                set (RB) i:51 \leftarrowinvalid
        if (MSR 
            for each process-scoped TLB entry in
            set (RB) i:51
                if entry_LPID=search_LPID
                    then TLB entry \leftarrow invalid
    if RIC=1 | RIC=2 then
        if MSR 
            if PRS=O then
                invalidate all partition-scoped
                page walk caching
            else
                invalidate all process-scoped
                page walk caching
        if (MSR 
                invalidate process-scoped page walk
                caching associated with partition
                search_LPID
    if RIC=2 then
    if MSR HV then
        if PRS=0 then
            invalidate all Partition Table caching
        else
                invalidate all Process Table caching
    if (MSR 
        invalidate Process Table caching associated
        with partition search_LPID
```

The operation performed by this instruction is based on the contents of registers RS and RB. The contents of these registers are shown below, where IS is $(\mathrm{RB})_{52: 53}$ and $L$ is $(R B)_{63}$.
RS:

|  | PID | $/ / /$ |
| :--- | :--- | :--- |
| 0 | 32 | 63 |

RB for $\mathrm{R}=1$ and $\mathrm{IS}=0 \mathrm{~b} 00$ ::

| RB for $\mathrm{R}=0$, $\mathrm{IS}=0 \mathrm{~b} 00$, and $\mathrm{L}=0$ :

| AVA | IS | B | AP | Os | L |
| :--- | :---: | :---: | :---: | :---: | :---: |
| 0 | 52 | 54 | 56 | 59 | 63 |

| RB for $R=0$, $I S=0 b 00$, and $L=1$ :

| AVA | LP | IS | B | AVAL | L |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 44 | 52 | 54 | 56 | 63 |

| RB for IS=0b01, Ob10, or Ob11:

| Os | SET | IS | Os |  |
| :---: | :---: | :---: | :---: | :---: |
| 0 | 40 | 52 | 54 | 63 |

LPIDR contains the partiion ID (LPID) of the partition for which the translation is being invalidated. $R S_{0: 31}$ contains a PID value. The supported values of $\mathrm{RS}_{0: 31}$ are the same as the PID values supported in PIDR.

The following forms are invalid.

- $\mathrm{PRS}=1, \mathrm{R}=0$, and $\mathrm{RIC} \neq 2$ (The only pro-cess-scoped HPT caching is of the Process Table.)
- RIC=1 and $\mathrm{R}=0$ (There is no Page Walk Cache for HPT translation.)
- RIC=3 (Cluster bombs are not supported for tlbiel.)

The following forms are treated as though the instruction form was invalid.

- RIC=1 and IS=0 (The architecture does not support shootdown of individual translations in the Page Walk Cache.)
- RIC=2 and IS=0 (RIC is for comprehensive invalidation that is not supported at the level of an individual page.)
- PRS=0 and IS=1 (Partition-scoped translations are not associated with processes.)
- $\mathrm{R}=0$ and $\mathrm{IS}=1$ (HPT translations are not associated with processes.)
The results of an attempt to invalidate a translation outside of quadrant 0 for Radix Tree translation ( $R=1$, RIC $=0$, $P R S=1, I S=0$, and $E A_{0: 1} \neq 0 b 00$ ) are boundedly undefined.


## IS field in RB contains Ob00

If RIC=0, this is a search for a single TLB entry. The following relationships must be true and tests and actions are performed to search for an HPT translation.

If the base page size specified by the PTE that was used to create the TLB entry to be invalidated is 4 $K B$, the $L$ field in register RB must contain 0 .
If the $L$ field in RB contains 0 , the base page size is 4 KB and $\mathrm{RB}_{56: 58}$ (AP - Actual Page size field) must be set to the SLBE size corresponding to the actual page size specified by the PTE that was used to create the TLB entry to be invalidated. Thus, $b$ is equal to 12 and $p$ is equal to $\log _{2}$ (actual page size specified by (RB) ${ }_{56: 58}$ ). The Abbreviated Virtual Address (AVA) field in register RB must contain bits 14:65 of the virtual address translated by the TLB entry to be invalidated. Variable i is equal to 51.

If the $L$ field in RB contains 1, the following rules apply.
■ The base page size and actual page size are specified in the LP field in register RB, where the relationship between (RB) 44:51 (LP - Large Page size selector field) and the base page size and actual page size is the same as the relationship between $P T E_{L P}$ and the base page size and actual page size, except for the "r" bits (see Section 5.7.9.1 on page 998 and Figure 31 on page 999). Thus, $b$ is equal to $\log _{2}$ (base page size specified by (RB) ${ }_{44: 51}$ ) and $p$ is equal to $\log _{2}$ (actual page size specified by (RB) ${ }_{44: 51}$ ). Specifically, (RB) ${ }_{44+\mathrm{c}: 51}$ must be equal to the contents of bits $\mathrm{c}: 7$ of the LP field of the PTE that was used to create the TLB entry to be invalidated, where $c$ is the maximum of 0 and ( $20-\mathrm{p}$ ).
■ Variable $i$ is the larger of ( $63-\mathrm{p}$ ) and the value that is the smaller of 43 and (63-b). (RB) $0_{0: i}$ must contain bits 14:(i+14) of the virtual address translated by the TLB to be invalidated. If $b>20, R_{64-b: 43}$ may contain any value and are ignored by the hardware.
■ If $b<20$, (RB) ${ }_{56: 75-b}$ must contain bits 58:77-b of the virtual address translated by the TLB to be invalidated, and other bits in (RB) $56: 62$ may contain any value and are ignored by the hardware.

- If $b \geq 20,(R B)_{56: 62}$ (AVAL - Abbreviated Virtual Address, Lower) may contain any value and are ignored by the hardware.
Let the segment size be equal to the segment size specified in (RB) ${ }_{54: 55}$ (B field). The contents of $\mathrm{RB}_{54: 55}$ must be the same as the contents of the $B$ field of the PTE that was used to create the TLB entry to be invalidated.
All TLB entries that have all of the following properties are made invalid on the thread executing the tlbiel instruction.
■ The entry translates a virtual address for which all the following are true.
- $\mathrm{VA}_{14: 14+\mathrm{i}}$ is equal to $(\mathrm{RB})_{0 ; i}$.
- $L=0$ or $b \geq 20$ or, if $L=1$ and $b<20$,
$\mathrm{VA}_{58: 77-\mathrm{b}}$ is equal to $(\mathrm{RB})_{56: 75-\mathrm{b}}$.
- The segment size of the entry is the same as the segment size specified in (RB) ${ }_{54: 55}$.
- Either of the following is true:
- The $L$ field in RB is 0 , the base page size of the entry is 4 KB , and the actual page size of the entry matches the actual page size specified in (RB) ${ }_{56: 58}$.
- The $L$ field in RB is 1, the base page size of the entry matches the base page size specified in (RB) ${ }_{44: 51}$, and the actual page size of the entry matches the actual page size specified in (RB) ${ }_{44: 51}$.
\| ■ $\operatorname{TLBE}_{\text {LPID }}=\operatorname{LPIDR}$ LPID.

The following relationships must be true and tests and actions are performed to search for a Radix Tree translation. For a partition-scoped invalidation, references to the effective address are understood to refer to the guest real address.

The page size is encoded in $\mathrm{RB}_{56: 58}$ (AP - Actual Page size field). Thus $p$ is equal to $\log _{2}$ ( page size specified by $\mathrm{RB}_{56: 58}$ ). The Effective Page Number (EPN) field in register RB must contain the bits 0:i of the effective address translated by the TLB entry to be invalidated. Variable $i$ is equal to $63-\mathrm{p}$.

The fields shown as zeros must be set to zero and are ignored by the hardware.
All TLB entries that have all of the following properties are made invalid on the thread executing the tlbiel instruction..

- The entry translates an effective address for which $E A_{0: i}$ is equal to $(R B)_{0: i}$.
■ The page size of the entry matches the page size specified in (RB) 56:58 $^{2}$.
- The entry has the appropriate scope (partition or process).
■ The process ID specified in RS matches the process ID in the TLB entry if not invalidating a partition-scoped translation.
- TLBE LPID matches the partiion ID of the partition for which the translation is to be invalidated.


## IS field in RB is non-zero

If RIC=0 or RIC=2, (RB) $)_{i: 51}$ (bits $i-40: 11$ of the SET field in (RB)) specify a set of TLB entries, where $i$ is an implementation-dependent value in the range 40:51. Each partition-scoped entry when PRS=0 and $\mathrm{MSR}_{\mathrm{HV}}=1$ or each process-scoped entry when PRS=1 in the set is invalidated if any of the following conditions are met for the entry.

- The IS field in RB contains 0b10, or $\mathrm{MSR}_{\mathrm{HV}}=0$ and the IS field contains 0b11, and TLBE $_{\text {LPID }}=$ LPIDR $_{\text {LPID. }}$
■ The IS field in RB contains 0b01, TLBE $_{\text {LPID }}=$ LPIDR $_{\text {LPID }}$, and TLBE $_{\text {PID }}=\mathrm{RS}_{0: 31}$.
- The IS field in RB contains 0b11 and $\mathrm{MSR}_{\mathrm{HV}}=1$.

How the TLB is divided into the $2^{52-i}$ sets is implemen-tation-dependent. The relationship of virtual addresses to these sets is also implementation-dependent. However, if, in an implementation, there can be multiple TLB entries for the same virtual address and same partition, then all these entries must be in a single set.
If $\mathrm{RIC}=1$ or $\mathrm{RIC}=2$, if the following conditions are met, the respective partition-scoped contents when $\mathrm{PRS}=0$ and $\mathrm{MSR}_{\mathrm{HV}}=1$ or process-scoped contents when $\mathrm{PRS}=1$ of the page walk cache are invalidated.

- If the IS field in RB contains $0 b 10$ or if IS contains Ob11 and $\mathrm{MSR}_{\mathrm{HV}}=0$, all properly-scoped page walk caching associated with partition LPDIR ${ }_{\text {LPID }}$ is invalidated.
- If the IS field in RB contains $0 b 11$ and $M S R_{H V}=1$, the entire properly-scoped page walk caching is invalidated.
- If the IS field in RB contains 0 001 (and PRS=1 and $\mathrm{R}=1$ ), all properly-scoped page walk caching associated with process $\mathrm{RS}_{0: 31}$ in partition LPIDR $_{\text {LPID }}$ is invalidated.
If $\mathrm{RIC}=2$, if the following conditions are met, the respective partition and Process Table caching are invalidated.
- If the IS field in RB contains $0 b 01$ and $P R S=1$, caching of Process Table Entries for process $\mathrm{RS}_{0: 31}$ in partition LPIDR ${ }_{\text {LPID }}$ is invalidated.
- If the IS field in RB contains $0 b 10$, MSR $_{H V}=1$, and PRS $=0$, caching of Partition Tables for partition LPIDR $_{\text {LPID }}$ is invalidated.
- If the is field in RB contains $0 b 10$ and PRS $=1$, caching of Process Tables for partition LPIDR is invalidated.
- if the IS field in RB contains $0 b 11, M S R_{H V}=1$, and PRS $=0$, all Partition Table caching is invalidated.
- if the IS field in RB contains 0b11, MSR ${ }_{H V}=1$, and PRS=1, all Process Table caching is invalidated.
- If the IS field in RB contains 0b11, MSR ${ }_{H V}=0$, and PRS=1, caching of Process Tables for partition LIDR $_{\text {LPID }}$ is invalidated.
When $\mathrm{i}>40, \mathrm{RB}_{40 \mathrm{i}:-1}$ may contain any value and are ignored by the hardware.


## For all IS values

Any implementation specific lookaside information that is based on any TLB entry that would be invalidated by this instruction will also be invalidated.
Depending on the variant of the instruction, $\mathrm{RS}_{32: 63}$, $\mathrm{RB}_{0: 39}, \mathrm{RB}_{59: 62}, \mathrm{RB}_{59: 63}, \mathrm{RB}_{54: 55}$, and $\mathrm{RB}_{54: 63}$ are the equivalent of reserved fields, should contain 0 s, and are ignored by the hardware.
Only TLB entries, page walk caching, and process and Segment Table caching on the thread executing the tlbiel instruction are affected.
$\mathrm{MSR}_{\text {SF }}$ must be 1 when this instruction is executed; otherwise the results are boundedly undefined.

If the value specified in $\mathrm{RS}_{0: 31}, \mathrm{RB}_{54: 55}, \mathrm{RB}_{56: 58}$, or $\mathrm{RB}_{44: 51}$, when it is needed to perform the specified operation, is not supported by the implementation, the instruction is treated as if the instruction form were invalid.
This instruction is privileged except when PRS=0 and $H R \| G R \neq 0 b 00$, making it hypervisor privileged.
See Section 5.10, "Translation Table Update Synchronization Requirements" on page 1043 for a description of other requirements associated with the use of this instruction.

## Special Registers Altered:

None

## Extended Mnemonics:

Extended mnemonic for tlbiel::
Extended:
tlbiel RB

## Equivalent to:

tlbiel RB,r0,0,0,0

## Programming Note

tlbie and t/biel serve as both basic and extended mnemonics. The Assembler will recognize a tlbie or tlbiel mnemonic with five operands as the basic form, and a tlbie with two operands or a tlbiel mnemonic with one operand as the extended form. In the extended form the RIC, PRS, and R operands, and for tlbiel the RS operand, are omitted and assumed to be 0 .

## Programming Note

The primary use of this instruction by hypervisor software is to invalidate TLB entries prior to reassigning a thread to a new logical partition.

For IS $\neq 0 \mathrm{ObOO}$, it is implementation-dependent whether ERAT entries are invalidated. If the tlbiel instruction is being executed due to a partition swap, an slbia instruction can be used to invalidate the pertinent ERAT entries. If the tlbiel instruction is being executed to invalidate TLB entries with parity or ECC errors, the fact that the corresponding ERAT entries are not invalidated is immaterial. If the tlbiel instruction is being executed to invalidate multiple matching TLB entries, the fact that the corresponding ERAT entries are not invalidated is immaterial for implementations that never create multiple matching ERAT entries.
The primary use of this instruction by operating system software is to invalidate TLB entries that were created by the hypervisor using an implemen-tation-specific hypervisor-managed TLB facility, if such a facility is provided.
tlbiel may be executed on a given thread even if the sequence tlbie - eieio - tlbsync - ptesync is concurrently being executed on another thread.
See also the Programming Notes with the description of the tlbie instruction.

## TLB Synchronize X-form

## tlbsync



The tlbsync instruction provides an ordering function for the effects of all tlbie instructions executed by the thread executing the tlbsync instruction, with respect to the memory barrier created by a subsequent ptesync instruction executed by the same thread. Executing a tlbsync instruction ensures that all of the following will occur.

- All TLB invalidations caused by tlbie instructions preceding the tlbsync instruction will have completed on any other thread before any data accesses caused by instructions following the ptesync instruction are performed with respect to that thread.
- All storage accesses by other threads for which the address was translated using the translations being invalidated, and all Reference and Change bit updates associated with address translations that were performed by other threads using the translations being invalidated, will have been performed with respect to the thread executing the ptesync instruction, to the extent required by the associated Memory Coherence Required attributes, before the ptesync instruction's memory barrier is created.
The operation performed by this instruction is ordered by the eieio (or sync or ptesync) instruction with respect to preceding tlbie instructions executed by the thread executing the tlbsync instruction. The operations caused by tlbie and tlbsync are ordered by eieio as a fourth set of operations, which is independent of the other three sets that eieio orders.
The tlbsync instruction may complete before operations caused by tlbie instructions preceding the t/bsync instruction have been performed.
This instruction is privileged except when $L_{\text {LPCR }}^{\text {GTSE }}=0$, making it hypervisor privileged.
See Section 5.10 for a description of other requirements associated with the use of this instruction.
Special Registers Altered:
None


## Programming Note

tlbsync should not be used to synchronize the completion of tlbiel.

### 5.10 Translation Table Update Synchronization Requirements

This section describes rules that software must follow when updating the Translation Tables, and includes suggested sequences of operations for some representative cases. The sequences required for other cases may be deduced from the sequences that are provided and from this accompanying description.

In the sequences of operations shown in the following subsections, the Page Table Entry is assumed to be for a virtual page for which the base page size is equal to the actual page size. If these page sizes are different, multiple tlbie instructions are needed, one for each PTE corresponding to the virtual page.
In the sequences of operations shown in the following | subsections, any alteration of a translation table entry that corresponds to a single line in the sequence is assumed to be done using a Store instruction for which the access is atomic. Appropriate modifications must be made to these sequences if this assumption is not satisfied (e.g., if a store doubleword operation is done using two Store Word instructions).
When storing to the bytes that contain the Reference and Change bits in a valid PTE or when setting the PTE invalid for the purpose of obtaining stable values for the Reference and Change bits, software in a Radix Tree translation environment must perform an atomic update of the PTE in order to interact correctly with hardware Reference and Change bit updates. In other circumstances, multithreaded software in a Radix Tree translation environment may use atomic updates and/ or some form of locking to ensure mutual exclusion on a PTE or other translation table entry it updates, subject to the limitation for a table entry with a valid bit that the entry must be made invalid (if initially valid) by the first store, and must be made valid by the final store to the entry (if it is made valid) for a lock-based sequence and that an atomic sequence (including those performed by hardware) must test for a valid entry before completing the sequence. The restriction is so that a valid table entry has a consistent state when modified using atomic update (including by hardware). Software in an HPT translation environment must use only lock-based sequences and obey the methodology just described. When the same type of sequence works for both types of translation, the HPT PTE is shown because it is more complex. In this description, and in references in subsequent subsections to "safe for multithreaded software," the safety is with respect to the risk of one thread overwriting another's update. There may also be concern for the creation of multiple matching translations, e.g. within a PTEG or pair of PTEGs. When the reservation granule is equal to or larger in size than the structure on which mutual exclusion must be ensured (e.g. PTE for Radix Tree translation but PTEG for HPT translation), multiple entries will also be prevented. (Secondary hash groups will generally not
be covered by the same reservation granule as primary hash groups.)

Updates (by software) to the tables are performed only when they are known to be required by the sequential execution model (see Section 5.5). Because address translation for instructions preceding a given Store instruction might cause an interrupt, and thereby prevent the corresponding store from being required by the sequential execution model, address translations for instructions preceding the Store instruction must be performed before the corresponding store is performed. As a result, an update to a translation table need not be preceded by a context synchronizing instruction.
All of the sequences require a context synchronizing operation after the sequence if the new contents of the | translation table are to be used for address translations associated with subsequent instructions.

As noted in the description of the Synchronize instruction in Section 4.6.3 of Book II, address translation associated with instructions which occur in program order subsequent to the Synchronize (and this includes the ptesync variant) may be performed prior to the completion of the Synchronize. To ensure that these instructions and data which may have been speculatively fetched are discarded, a context synchronizing operation is required.

## - Programming Note

In many cases this context synchronization will occur naturally; for example, if the sequence is executed within an interrupt handler the rfid, rfscv, or hrfid instruction that returns from the interrupt handler may provide the required context synchronization.
| Translation table entries must not be changed in a manner that causes an implicit branch.

### 5.10.1 Translation Table Updates

TLBs are non-coherent caches of the HTABs and Radix Trees. TLB entries must be invalidated explicitly with one of the TLB Invalidate instructions. SLBs are non-coherent caches of the Segment Tables, SLB entries must be invalidated explicitly with one of the SLB Invalidate instructions. Page Walk Caches are non-coherent caches of the intermediate steps in Radix Tree translation. Non-coherent caching of the partition and Process Tables is permitted. Provision has been made for the use of the TLB Invalidate instructions to manage the types of caching described in the preceding two sentences at a PID or LPID granularity.

Unsynchronized lookups in the Page, Segment, and when HRIIGR=0b00, Process Tables continue even
while they are being modified. (For Partition Table Entries, and for Process Table Entries when HR\|GR $\neq 0 b 00$, the process or partition affected must be inactive because the entries do not have valid bits.) With the exceptions previously identified for Segment Table walks (see Section 5.9.3, "Lookaside Buffer Management"), any thread, including a thread on which software is modifying any of the set of tables described in the first sentence, may look in those tables at any time in an attempt to translate an address. When modifying an entry in any of the former set of tables, software must ensure that the table entry's V bit is 0 if the table entry does not correctly specify its portion of the translation (e.g., if the RPN field is not correct for the current AVA field).
| For HPT translation, updates of Reference and Change bits by the hardware are not synchronized with the accesses that cause the updates. When modifying doubleword 1 of a PTE, software must take care to avoid overwriting a hardware update of these bits and to avoid having the value written by a Store instruction overwritten by a hardware update.

The most basic sequence that will achieve proper system synchronization for PTE updates is the following.
tlbie instruction(s) specifying the same LPID operand value

## eieio

tlbsync
ptesync
Other instructions may be interleaved among these instructions. Operating system and hypervisor software that updates Page Table Entries should use this sequence.
Operating systems and nested hypervisors are exposed to being interrupted during this sequence. The interrupting hypervisor is responsible for completing the sequence above. In general this will require the hypervisor to include the following sequence in an interrupt handler.

```
eieio
tlbsync
ptesync
```

This sequence itself may be interrupted by a higher level hypervisor. When returning to the interrupted software, the original sequence will be completed. Hardware must tolerate the result of nested interleaving of these sequences. tlbie and tlbsync instructions should only be used as part of these sequences.
The corresponding sequences for Segment Table updates use slbieg in place of tlbie and slbsync in place of tlbsync. Similarly slbieg and slbsync should only be used as part of these sequences. In circumstances where a hypervisor may be interrupting either a PTE update or a Segment Table update, it must include both tlbsync and slbsync in its completing sequence, in either order. Hardware must tolerate the result of nested interleaving of these additional sequences.

The PTE sequences are also used to synchronize updates to partition and Process Table Entries.
On systems consisting of only a single-threaded processor, the eieio and tlbsync or slbsync instructions can be omitted.

The following subsections illustrate sequences that must be used for translation table updates to tables that are subject to concurrent use by hardware (i.e. that have valid bits in their entries). For Partition Table Entries and for Process Table Entries that do not have valid bits, simpler sequences consisting of just the preceding sequences, perhaps with mutual exclusion if the update processes are multithreaded, is sufficient.

## Programming Note

The eieio instruction prevents the reordering of the preceding tlbie or slbieg instructions with respect to the subsequent tlbsync or slbsync instruction. The tlbsync or slbsync instruction and the subsequent ptesync instruction together ensure that all storage accesses for which the address was translated using the translations being invalidated (by the tlbie or slbieg instructions), and all Reference and Change bit updates associated with address translations that were performed using the translations being invalidated, will be performed with respect to any thread or mechanism, to the extent required by the associated Memory Coherence Required attributes, before any data accesses caused by instructions following the ptesync instruction are performed with respect to that thread or mechanism.

For Page Table update sequences that mark the PTE invalid (see Section 5.10.1.2, "Modifying a Translation Table Entry"), Reference and Change bit updates cease when the sequence is complete. When the PTE is marked invalid using an atomic update and the Store Conditional setting the entry invalid is successful, the Reference and Change bits obtained by the corresponding Load And Reserve instruction are stable/final values.

The sequences of operations shown in the following subsections assume a multi-threaded environment. In an environment consisting of only a single-threaded processor, the tlbsync or slbsync and the eieio that separates the tlbie or slbieg from the tlbsync or slbsync can be omitted. In a multi-threaded environment, when tlbiel or slbie is used instead of tlbie or slbieg in a Page or Segment Table update, the synchronization requirements are the same as when tlbie or slbieg is used in an environment consisting of only a sin-gle-threaded processor.

## - Programming Note

For all of the sequences shown in the following subsections, if it is necessary to communicate completion of the sequence to software running on another thread, the ptesync instruction at the end of the sequence should be followed by a Store instruction that stores a chosen value to some chosen storage location X. The memory barrier created by the ptesync instruction ensures that if a Load instruction executed by another thread returns the chosen value from location X, all subsequent searches of the Page or Segment Table by the other thread, that implicitly load from the PTE or STE specified by the sequence's stores, will obtain the values stored (or values stored subsequently). The Load instruction that returns the chosen value should be followed by a context synchronizing instruction in order to ensure that all instructions following the context synchronizing instruction will be fetched and executed using the values stored by the sequence (or values stored subsequently). (These instructions may have been fetched or executed out-of-order using the old contents of the PTE or STE.)

This Note assumes that the Page or Segment Table and location X are in storage that is Memory Coherence Required.

### 5.10.1.1 Adding a Page Table Entry

This is the simplest Page Table case. The V bit of the old entry is assumed to be 0 . The following sequence can be used to create a PTE, maintain a consistent state, and ensure that a subsequent reference to the virtual address translated by the new entry will use the correct real address and associated attributes. A single quadword store would avoid the need for the eieio. A similar sequence may be used to add a new Segment Table Entry. Mutual exclusion with respect to other software threads may be required, but there is no concern for interaction with hardware updates because the entry is invalid until the last store in the sequence.
| PTE $\mathrm{p}_{\mathrm{pp}}$ key B ARPN LP key R C WImG N pp $\leftarrow$ new values
eieio /* order 1st update before 2nd */
\| $\mathrm{PTE}_{\text {AVA }} S W \mathrm{LH} V \leftarrow$ new values (V=1)
ptesync /* order updates before next
Page Table search and before next data access */

## | 5.10.1.2 Modifying a Translation Table Entry

| General Case (PTE)
If a valid entry is to be modified and the translation instantiated by the entry being modified is to be invali-
I dated, the sequences below can be used to modify the

PTE, maintain a consistent state, ensure that the translation instantiated by the old entry is no longer available, and ensure that a subsequent reference to the virtual address translated by the new entry will use the correct real address and associated attributes.

The following sequence is for Radix Tree translation. It interacts correctly with hardware atomic updates to return stable Reference and Change bit values for the old translation and is safe for multitheaded software. If the purpose of the sequence is mainly to collect Reference and Change bit values, the part of the sequence beginning with tlbie may be deferred and performed as a bulk invalidation (e.g. for a range of storage or an entire process) after collecting values for a plurality of pages. A similar seqence (i.e. using Load And Reserve and Store Conditional instructions) can be used to update a Segment Table Entry but cannot be used to update an HPT PTE because it will not interact correctly with non-atomic hardware Reference and Change bit updates.

```
r6\leftarrowPTE \ L SW RPN R C Att EAA
r4\leftarrowaddr (pte)
loop:
    lqarx r2,0,r4
    if V=0 abort, else /* to interact with locking */
    stqcx r6,0,r4
    bne- loop
ptesync /* order update before tlbie and
                                    before next Page Table search */
tlbie(old_EA 0:63-b,old_AP,old_PID,
            old_LPID)
eieio /* order tlbie before tlbsync */
tlbsync /* order tlbie before ptesync */
ptesync /*complete the sequence, stores ordered
    /*by first ptesync
```

The corresponding sequence for HPT translation is the following. (The sequence is equivalent to deleting the PTE and then adding a new one.) Mutual exclusion with respect to other software threads may be required. The Reference and Change bit values will not be stable until the entire sequence is completed.

```
PTE 
ptesync /* order update before tlbie and
    before next Page Table search */
tlbie(old_B,old_VA14:77-b,old_L,old_LP,old_AP,
    old_LPID)
    /*invalidate old translation*/
eieio /* order tlbie before tlbsync */
tlbsync /* order tlbie before ptesync */
ptesync /* order tlbie, tlbsync and 1st
    update before 2nd update */
PTE ARPN,LP,AC,R,C,WIMG,N,PP}\leftarrow new values
eieio /* order 2nd update before 3rd */
PTE B,AVA, SW, L, H,V }\leftarrow new values (V=1
ptesync /* order 2nd and 3rd updates before
        next Page Table search and
        before next data access */
```


## General Case(STE)

If a valid entry is to be modified and the translation instantiated by the entry being modified is to be invali-
dated, the following sequence can be used to modify the STE, maintain a consistent state, ensure that the translation instantiated by the old entry is no longer available, and ensure that a subsequent reference to the effective address translated by the new entry will use the correct virtual address and associated attributes. (The sequence is much like the general case for a change to an HPT PTE, and is equivalent to deleting the STE and then adding a new one.) Mutual exclusion with respect to other software threads may be required. A similar sequence (except using t/bie with RIC=2 and tlbsync) may be used to modify HRIIGR=0b00 Process Table Entries.

```
STE 
ptesync /* order update before slbieg and
        before next Segment Table search */
slbieg(old_B,old_ESID,old_C,old_TA)
        /*invalidate old translation*/
    eieio /* order slbieg before slbsync */
    slbsync /* order slbieg before ptesync */
ptesync /* order slbieg, slbsync and 1st
        update before 2nd update */
    /* deletion sequence ends here */
    STE VSID, Ks, Kp, N, L, C, LP, SW }\leftarrow new value
    eieio /* order 2nd update before 3rd */
```

\|
I
Resetting the Reference Bit (PTE)

If the only change being made to a valid entry is to set the Reference bit to 0 , a simpler sequence suffices because the Reference bit need not be maintained exactly. The byte store is exposed to overwriting another change being performed by multithreaded software, so mutual exclusion may be required.

```
oldR \leftarrow PTE R /* get old R */
if oldR = 1 then
    PTE R < % /* store byte (R=0, other bits
    tlbie(B, VA 14:77-b,L,LP,AP,LPID) /* invalidate
    eieio /* order tlbie before tlbsync */
    tlbsync /* order tlbie before ptesync */
    ptesync /* order tlbie, tlbsync, and update
            before next Page Table search
            and before next data access */
```


## Modifying the SW field (PTE)

If the only change being made to a valid entry is to modify the SW field, the following sequence suffices, because the SW field is not used by the hardware (i.e. is not cached in the TLB and has no effect on hardware behavior).

[^14]```
stdcx. PTE_dwd_0 \leftarrow r1 /* store dwd 0 of PTE
        if still reserved (new SW value, other
        fields unchanged) */
bne- loop /* loop if lost reservation */
```

A Ibarx/stbcx., Iharx/sthcx., or Iwarx/stwcx. pair (specifying the low-order byte, halfword, or word respectively of doubleword 0 of the PTE) can be used instead of the Idarx /stdcx. pair shown above for HPT translation. The split SW field in the radix PTE cannot be updated with a single smaller atomic update. This sequence interacts correctly with hardware updates and is safe for multithreaded software. A similar sequence (including the possibility of using a smaller atomic update) can be used to update a Segment Table Entry.

## Modifying the Effective Address (STE)

If the effective address translated by a valid STE is to be modified and the new effective address hashes to the same STEG as does the old effective address, the following sequence can be used to modify the STE, maintain a consistent state, ensure that the translation instantiated by the old entry is no longer available, and ensure that a subsequent reference to the effective address translated by the new entry will use the correct virtual address and associated attributes. Mutual exclusion with respect to other software threads may be required. The corresponding change of the virtual address in the PTE for HPT translation can be performed using a similar sequence, interacting correctly with hardware table updates, as long as the second doubleword of the PTE is not stored.

[^15]
## Chapter 6. Interrupts

### 6.1 Overview

The Power ISA provides an interrupt mechanism to allow the thread to change state as a result of external signals, errors, or unusual conditions arising in the execution of instructions.

System Reset and Machine Check interrupts are not ordered. All other interrupts are ordered such that only one interrupt is reported, and when it is processed (taken) no program state is lost. Since Save/Restore Registers SRRO and SRR1 are serially reusable resources used by most interrupts, program state may be lost when an unordered interrupt is taken.

### 6.2 Interrupt Registers

### 6.2.1 Machine Status Save/ Restore Registers

When various interrupts occur, the state of the machine is saved in the Machine Status Save/Restore registers (SRR0 and SRR1). Section 6.5 describes which registers are altered by each interrupt.


Figure 51. Save/Restore Registers
SRR1 bits may be treated as reserved in a given implementation if they correspond to MSR bits that are reserved or are treated as reserved in that implementation and, for SRR1 bits in the range 33:36, 42:43, and 45:47, they are specified as being set either to 0 or to an undefined value for all interrupts that set SRR1 (including implementation-dependent setting, e.g. by the Machine Check interrupt or by implementation-specific interrupts). SRR1 ${ }_{44}$ cannot be treated as reserved, regardless of how it is set by interrupts, because it is used by software, as described in a Programming Note
near the end of Section 6.5.9, "Program Interrupt" on page 1071.

### 6.2.2 Hypervisor Machine Status Save/Restore Registers

When various interrupts occur, the state of the machine is saved in the Hypervisor Machine Status Save/ Restore registers (HSRR0 and HSRR1). Section 6.5 describes which registers are altered by each interrupt.


Figure 52. Hypervisor Save/Restore Registers
HSRR1 bits may be treated as reserved in a given implementation if they correspond to MSR bits that are reserved or are treated as reserved in that implementation and, for HSRR1 bits in the range 33:36 and 42:47, they are specified as being set either to 0 or to an undefined value for all interrupts that set HSRR1 (including implementation-dependent setting, e.g. by implementa-tion-specific interrupts).
The HSRR0 and HSRR1 are hypervisor resources; see Chapter 2.

## Programming Note

Execution of some instructions, and fetching instructions when $\mathrm{MSR}_{\mathrm{IR}}=1$, may have the side effect of modifying HSRR0 and HSRR1; see Section 6.4.4.

### 6.2.3 Access Segment Descriptor Register

The DAR, HDAR, SRR0, and HSRR0 generally provide the EA for storage exceptions. For hypervisor storage interrupts, additional information is often necessary to enable the hypervisor to handle the interrupt. This information is provided in a 64b SPR called the Access

Segment Descriptor Register (ASDR). When nested translation is taking place, the ASDR will generally provide the guest real address down to bit 51. (The smallest supported page size is 4 k .) When using paravirtualized HPT translation or for HV=1 accesses when $H R=0$, information from the segment descriptor that was used to perform the effective to virtual translation is provided in the ASDR. For exceptions that take place when translating the address of the process table entry or segment table entry group, only the VSID will be provided, because those addresses are specified as virtual addresses and the rest of the segment descriptor is implied. Some instances of the Machine Check interrupt may require the ASDR to be set similarly to how it is set for the hypervisor storage interrupts. The ASDR is set independent of the value of UPRT for the partition that is running.

| B |  | VSID | $\mathrm{K}_{\mathrm{s}} \mathrm{K}_{\mathrm{p}} \mathrm{NLC}$ | $/$ | LP |
| :--- | :--- | :--- | :--- | :--- | :--- |

Figure 53. Access Segment Descriptor Register format for a Segment Descriptor

| $/$ | GRA | 0 |  |
| :---: | :---: | :---: | :---: |
| 0 | 2 | 52 | 63 |

Figure 54. Access Segment Descriptor Register format for a Guest Real Address

### 6.2.4 Data Address Register

The Data Address Register (DAR) is a 64-bit register that is set by the Machine Check, Data Storage, Data Segment, and Alignment interrupts; see Sections 6.5.2, 6.5.3, 6.5.4, and 6.5.8. In general, when one of these interrupts occurs the DAR is set to an effective address associated with the storage access that caused the interrupt, with the high-order 32 bits of the DAR set to 0 if the interrupt occurs in 32-bit mode.


Figure 55. Data Address Register

### 6.2.5 Hypervisor Data Address Register

The Hypervisor Data Address Register (HDAR) is a 64-bit register that is set by the Hypervisor Data Storage Interrupt; see Section 6.5.16. In general, when this interrupt occurs, the HDAR is set to an effective address associated with the storage access that
caused the interrupt, with the high-order 32 bits of the HDAR set to 0 if the interrupt occurs in 32-bit mode.


Figure 56. Hypervisor Data Address Register

### 6.2.6 Data Storage Interrupt Status Register

The Data Storage Interrupt Status Register (DSISR) is a 32-bit register that is set by the Machine Check, Data I Storage, and Data Segment interrupts; see Sections 6.5.2, 6.5.3, and 6.5.4.


Figure 57. Data Storage Interrupt Status Register DSISR bits may be treated as reserved in a given implementation if they are specified as being set either to 0 or to an undefined value for all interrupts that set the DSISR.

### 6.2.7 Hypervisor Data Storage Interrupt Status Register

The Hypervisor Data Storage Interrupt Status Register (HDSISR) is a 32-bit register that is set by the Hypervisor Data Storage interrupt. In general, when one of these interrupts occurs the HDSISR is set to indicate the cause of the interrupt.


Figure 58. Hypervisor Data Storage Interrupt Status Register

### 6.2.8 Hypervisor Emulation Instruction Register

The Hypervisor Emulation Instruction Register (HEIR) is a 32-bit register that is set by the Hypervisor Emulation Assistance interrupt; see Section 6.5.18. The image of the instruction that caused the interrupt is loaded into the register.


Figure 59. Hypervisor Emulation Instruction Register

### 6.2.9 Hypervisor Maintenance Exception Register

Each bit in the Hypervisor Maintenance Exception Register (HMER) is associated with one or more causes of the Hypervisor Maintenance exception, and is set when the associated exception(s) occur. If the corresponding bit in the Hypervisor Maintenance Exception Enable Register (HMEER) is set, a Hypervisor Maintenance Interrupt (HMI) may occur. If the thread is in a power-saving mode when the interrupt would have occurred, the thread will exit the power-saving mode; see Section 6.5.19 and Section 3.3.2.


Figure 60. Hypervisor Maintenance Exception Register

The contents of the HMER are as follows:
$0 \quad$ Set to 1 for a Malfunction Alert.
1 Set to 1 when performance is degraded for thermal reasons.
2 Set to 1 when thread recovery is invoked.
Others Implementation-specific.
When the mtspr instruction is executed with the HMER as the encoded Special Purpose Register, the contents of register RS are ANDed with the contents of the HMER and the result is placed into the HMER.

The exception bits in the HMER are sticky; that is, once set to 1 they remain set to 1 until they are set to 0 by an mthmer instruction.

## Programming Note

An access to the HMER is likely to be very slow. Software should access it sparingly.

### 6.2.10 Hypervisor Maintenance Exception Enable Register

The Hypervisor Maintenance Exception Enable Register (HMEER) is a 64-bit register in which each bit enables the corresponding exception in the HMER to cause the Hypervisor Maintenance interrupt, potentially causing exit from power-saving mode; see Section 6.5.19 and Section 3.3.2.


Figure 61. Hypervisor Maintenance Exception Enable Register

### 6.2.11 Facility Status and Control Register

The Facility Status and Control Register (FSCR) controls the availability of various facilities in problem state and indicates the cause of a Facility Unavailable interrupt.
When the FSCR makes a facility unavailable, attempted usage of the facility in problem state is treated as follows:

- Execution of an instruction causes a Facility Unavailable interrupt.
- Access of an SPR using mfspr/mtspr causes a Facility Unavailable interrupt
- rfebb, rfid, rfscv, hrfid and mtmsr[d] instructions have the same effect on bits in system registers as they would if the bits were available.


## Programming Note

The FSCR does not prevent rfebb instructions from attempting to set bits in System Registers that the FSCR makes unavailable. Thus changes to BES$\mathrm{CR}_{\text {TS }}$ made by the operating system have the potential to result in an illegal transaction state transition when rfebb is subsequently executed in problem state, resulting in the occurrence of a TM Bad Thing type Program interrupt.

The MSR can also make the Transactional Memory facility unavailable in any privilege state, and MMCRO can make various components of the Performance Monitor unavailable when accessed in problem state. An access to one of these facilities when it is unavailable causes a Facility Unavailable interrupt.
When the PCR makes a facility unavailable in problem state, the facility is treated as not implemented in problem state; any Facility Unavailable interrupt that would occur if the facility were not made unavailble by the PCR does not occur.
When a Facility Unavailable interrupt occurs, the unavailable facility that was accessed is indicated in the most-significant byte of the FSCR.

| IC | Facility Control |  |
| :--- | :--- | :--- |
| 0 | 8 | 63 |

Figure 62. Facility Status and Control Register
The contents of the FSCR are specified below.

## Value Meaning

0:7 Interruption Cause (IC)
When a Facility Unavailable interrupt occurs, the IC field contains a binary number indicating the facility for which access was
attempted. The values and their meanings are specified below.

02 Access to the DSCR at SPR 3
03 Access to a Performance Monitor SPR in group A or B when $\mathrm{MMCRO}_{\text {PMCC }}$ is set to a value for which the access results in a Facility Unavailable interrupt. (See the definition of MMCRO ${ }_{\text {PMCC }}$ in Section 9.4.4.)

04 Execution of a BHRB Instruction
05 Access to a Transactional Memory SPR or execution of a Transactional Memory Instruction
06 Reserved
07 Access to an Event-Based Branch SPR or execution of an Event-Based Branch instruction
08 Access to the Target Address Register
OA Access to the msgsndp or msgclrp instructions, the TIR or the DPDES Register
OB Access to the Load Monitored Region Register or Load Monitored Section Enable Register.
OC Execution of scv
All other values are reserved.
Facility Enable (FE)
The FE field controls the availability of various facilities in problem state as specified below.

0 The msgsndp and msgclrp instructions and the TIR and DPDES registers are not available in privileged non-hypervisor state.
1 The msgsndp and msgcIrp instructions and the TIR and DPDES registers are available in privileged non-hypervisor state unless made unavailable by another register.

## Reserved

## Target Address Register (TAR)

0 The TAR and bctar instruction are not available in problem state.
1 The TAR and bctar instruction are available in problem state unless made unavailable by another register.

## Event-Based Branch Facility (EBB)

0 The Event-Based Branch facility SPRs and instructions are not available in problem state, and event-based exceptions and branches do not occur.
1 The Event-Based Branch facility SPRs and instructions (see Chapter 7 of Book II) are available in problem state unless made unavailable by another register, and event-based exceptions and branches are allowed to occur if enabled by other registers.
Reserved

## Programming Note

$\mathrm{HFSCR}_{58: 60}$ are used to control the availability of Transactional Memory, the Performance Monitor, and the BHRB in problem and privileged non-hypervisor states. $\mathrm{FSCR}_{58: 60}$ are reserved since the availability of Transactional Memory is controlled by the MSR, and the availability of the Performance Monitor and BHRB is controlled by MMCRO.

61 Data Stream Control Register at SPR 3 (DSCR)
0 SPR 3 is not available in problem state.
1 SPR 3 is available in problem state unless made unavailable by another register.

62:63 Reserved

## Programming Note

When an OS has set the FSCR such that a facility is unavailable, the OS should either emulate the facility when it is accessed or provide an application interface that requires the application to request use of the facility before it accesses the facility.

### 6.2.12 Hypervisor Facility Status and Control Register

The Hypervisor Facility Status and Control Register (HFSCR) controls the availability of various facilities in problem and privileged non-hypervisor states, and indicates the cause of a Hypervisor Facility Unavailable interrupt.

When the HFSCR makes a facility unavailable, attempted usage of the facility in problem or privileged non-hypervisor states is treated as follows:

- Execution of an instruction causes a Hypervisor Facility Unavailable interrupt.
- Access of an SPR using mfspr/mtspr causes a Hypervisor Facility Unavailable interrupt
| - rfebb, rfid, rfscv, hrfid and mtmsr[d] instructions have the same effect on bits in system registers as they would if the bits were available.


## Programming Note

Because the HFSCR does not prevent mtspr,
ting bits in system registers that the HFSCR will make unavailable after a transition to a lower privilege state, these instructions may cause interrupts in a variety of unexpected ways. For example, consider a hypervisor that sets HSRR1 such that hrfid returns to a lower privilege state with MSR[TS] nonzero. A TM Bad Thing type Program interrupt will result, despite that TM is made unavailable by the HFSCR.

Similarly, the HFSCR does not prevent rfebb instructions from attempting to set bits in System Registers that the HFSCR makes unavailable. Thus changes to BESCR ${ }_{\text {TS }}$ made by the hypervisor have the potential to result in an illegal transaction state transition when rfebb is subsequently executed in problem or privileged state, resulting in the occurrence of a TM Bad Thing type Program interrupt.

When the PCR makes a facility unavailable in problem state, the facility is treated as not implemented in problem state; any Hypervisor Facility Unavailable interrupt that would occur if the facility were not made unavailble by the PCR does not occur as a result of problem state access. See Section 2.5 for additional information.)

When a Hypervisor Facility Unavailable interrupt occurs, the facility that was accessed is indicated in the most-significant byte of the HFSCR.

| IC | Facility Control |  |
| :--- | :--- | :--- |
| 0 | 8 | 63 |

Figure 63. Hypervisor Facility Status and Control Register

The contents of the HFSCR are specified below.

## Value Meaning

0:7 Interruption Cause (IC)
When a Hypervisor Facility Unavailable interrupt occurs, the IC field contains a binary number indicating the access that was attempted. The values and their meanings are specified below.
00 Access to a Floating Point register or execution of a Floating Point instruction
01 Access to a Vector or VSX register or execution of a Vector or VSX instruction
02 Access to the DSCR at SPRs 3 or 17
03 Read or write access of a Performance Monitor SPR in group A, or read access of a Performance Monitor SPR in group B. (See Section 9.4.1 for a definition of groups A and B.)
04 Execution of a BHRB Instruction
05 Access to a Transactional Memory SPR or execution of a Transactional Memory instruction
06 Reserved
07 Access to an Event-Based Branch SPR or execution of an Event-Based Branch instruction
08 Access to the Target Address Register
09 Access to the stop instruction in privileged non-hypervisor state when one or more of the following conditions exist.
$\mathrm{PSSCR}_{E C}=1$
PSSCR $_{\text {ESL }}=1$
PSSCR $_{\text {MTL }}>$ PSSCR $_{\text {PSLL }}$
PSSCR $_{\text {RL }}>$ PSSCR $_{\text {PSLL }}$
All other values are reserved.
8:63 Facility Enable (FE)
The FE field controls the availability of various facilities in problem and privileged non-hypervisor states as specified below.
8:54 Reserved

## Programming Note

There is no bit in this register controlling the availability of the stop instruction because the availability of stop in privileged non-hypervisor state is controlled by the PSSCR. See Section 3.2.3.

There is no bit in this register controlling the availability of the Load Monitored Region Register, the Load Monitored Section Enable Register, and the Idmx instruction because no need for the hypervisor to control this availability has been identified.

Target Address Register (TAR)
0 The TAR and bctar instruction are not available in problem and privileged non-hypervisor state.
1 The TAR and bctar instruction are available in problem and privileged states unless made unavailable by another register.

## Event-Based Branch Facility (EBB)

0 The Event-Based Branch facility SPRs and instructions are not available in problem and privileged non-hypervisor states, and event-based exceptions and branches do not occur.
1 The Event-Based Branch facility SPRs and instructions are available in problem and privileged states unless made unavailable by another register, and event-based exceptions and branches are allowed to occur if enabled by other bits.
Reserved
Transactional Memory Facility (TM)
0 The Transactional Memory Facility SPRs and instructions are not available in problem and privileged non-hypervisor states.
1 The Transactional Memory Facility SPRs and instructions are available in problem and privileged states unless made unavailable by another register.
BHRB Instructions (BHRB)
0 The BHRB instructions (clrbhrb, mfbhrbe) are not available in problem and privileged non-hypervisor states.
1 The BHRB instructions (clrbhrb, mfbhrbe) are available in problem and privileged states unless made unavailable by another register.
Performance Monitor Facility SPRs (PM)
0 Read and write operations of Performance Monitor SPRs in group A and read opera-
tions of Performance Monitor SPRs in group $B$ are not available in problem and privileged non-hypervisor states; read and write operations to privileged Performance Monitor registers (SPRs 784-792, 795-798) are not available in privileged non-hypervisor state. (See Section 9.4.1 for a definition of groups A and B.) Performance Monitor exceptions do not cause Performance Monitor interrupts to occur when the thread is in problem or privileged states.
1 Read and write operations of Performance Monitor SPRs in group A and read operations of Performance Monitor SPRs in group $B$ are available in problem and privileged states unless made unavailable by another register; read and write operations to privileged Performance Monitor registers (SPRs 784-792, 795-798) are available in privileged state; Performance Monitor interrupts to occur if $\mathrm{MSR}_{\text {EE }}=1$ and MMCRO $_{\text {EBE }}=0$. See Section 9.2 of Book III for additional information

## Data Stream Control Register (DSCR)

0 SPR 3 is not available in problem or privileged non-hypervisor states and SPR 17 is not available in privileged non-hypervisor state.
1 SPR 3 is available in problem and privileged states and SPR 17 is available in privileged state unless made unavailable by another register.
62 Vector and VSX Facilities (VECVSX)
0 The facilities whose availability is controlled by either $\mathrm{MSR}_{\text {VEC }}$ or $\mathrm{MSR}_{\text {VSX }}$ are not available in problem and privileged non-hypervisor states.
1 The facilities whose availability is controled by either $\mathrm{MSR}_{\text {VEC }}$ or $\mathrm{MSR}_{\mathrm{VSX}}$ are available in problem and privileged states unless made unavailable by another register.
Floating Point Facility (FP)
0 The facilities whose availability is controlled by $\mathrm{MSR}_{\text {FP }}$ are not available in problem and privileged non-hypervisor states.
1 The facilities whose availability is controlled by MSR $_{\text {FP }}$ are available in problem and privileged states unless made unavailable by another register.

## Programming Note

The FSCR can be used to determine whether a particular facility is being used by an application, and the HFSCR can be used to determine whether a particular facility is being used by either an application or by an operating system. This is done by disabling the facility initially, and enabling it in the interrupt handler upon first usage. The information about the usage of a particular facility can be used to determine whether that facility's state must be saved and restored when changing program context.

## Programming Note

The following tables summarize the interrupts that occur as a result of accessing the non-privileged Performance Monitor registers in problem state when MMCRO ${ }_{\text {PMCC }}$, PCR, and HFSCR are set to various values. (Accesses to privileged Performance Monitor SPRs (SPRs 784-792, 795-798) in problem state result in Privileged Instruction Type Program interrupts.)


Notes:

1. Terminology:

FU: Facility Unavailable interrupt
HE: Hypervisor Emulation Assistance interrupt
HU: Hypervisor Facility Unavailable interrupt
2. This SPR is read-only, and cannot be written in any privilege state. (See the mtspr instruction description in Section 4.4.5 for additional information.) FU or HU interrupts do not occur regardless of the value of MMCRO PMCC or HFSCR ${ }_{\text {PM }}$.
3. When the PCR indicates a version of the architecture prior to $V 2.07$, this SPR is treated as not implemented in problem state; no FU or HU interrupts occur regardless of the value of MMCRO $\mathrm{PMCC}_{\text {P }}$ or $H_{F S C R}$ PM.
4. An HU interrupt occurs if $\mathrm{HFSCR}_{\mathrm{PM}}=0$ when this SPR is accessed in either problem state or privileged non-hypervisor state.

## Programming Note

When an MSR bit makes a facility unavailable, the facility is made unavailable in all privilege states. Examples of this include the Floating Point, Vector, and VSX facilities. The FSCR and HFSCR affect the availability of facilities only in privilege states that are lower than the privilege of the register (FSCR or HFSCR).

### 6.3 Interrupt Synchronization

I When an interrupt occurs, in general SRRO or HSRRO is set to point to an instruction such that all preceding instructions have completed execution, no subsequent instruction has begun execution, and the instruction addressed by SRRO or HSRRO may or may not have completed execution, depending on the interrupt type. The only exception is that if an mtspr sequence started by mtgsr is active when the interrupt occurs, some of the sequence's mtsprs beyond the instruction pointed to by SRRO or HSRRO may have been executed; see Chapter 11.
With the exception of System Reset and Machine Check interrupts, all interrupts are context synchronizing as defined in Section 1.5.1. System Reset and Machine Check interrupts are context synchronizing if they are recoverable (i.e., if bit 62 of SRR1 is set to 1 by the interrupt). If a System Reset or Machine Check interrupt is not recoverable (i.e., if bit 62 of SRR1 is set to 0 by the interrupt), it acts like a context synchronizing operation with respect to subsequent instructions. That is, a non-recoverable System Reset or Machine Check interrupt need not satisfy items 1 through 3 of Section 1.5.1, but does satisfy items 4 and 5 .

## Programming Note

While the behavior of correct mtgsr-initiated mtspr sequences described above does not directly violate a careful interpretation of the requirements of context synchronization, it may violate programmer expectations. For example, speculative execution is permitted

### 6.4 Interrupt Classes

Interrupts are classified by whether they are directly caused by the execution of an instruction or are caused by some other system exception. Those that are "sys-tem-caused" are:

- System Reset
- Machine Check
- External
- Decrementer
- Directed Privileged Doorbell
- Hypervisor Decrementer
- Hypervisor Maintenance

I - Hypervisor Virtualization

- Directed Hypervisor Doorbell
- Performance Monitor

External, Decrementer, Hypervisor Decrementer, Directed Privileged Doorbell, Directed Hypervisor Doorbell, Hypervisor Maintenance, and Hypervisor Virtualization interrupts are maskable interrupts. Therefore, software may delay the generation of these
interrupts. System Reset and Machine Check interrupts are not maskable.
"Instruction-caused" interrupts are further divided into two classes, precise and imprecise.

### 6.4.1 Precise Interrupt

I All instruction-caused interrupts other than the Imprecise Mode Floating-Point Enabled Exception type Program interrupt are precise, except that if an mtspr sequence started by $\boldsymbol{m t g s r}$ is active when the interrupt occurs, some of the sequence's mtsprs beyond the interrupt point may have been executed; see Chapter 11. Statements elsewhere in Book III-S that a given interrupt is precise do not preclude the mtspr case just described.
When the fetching or execution of an instruction causes a precise interrupt, the following conditions exist at the interrupt point.

1. SRRO addresses either the instruction causing the exception or the immediately following instruction. Which instruction is addressed can be determined from the interrupt type and status bits.
2. An interrupt is generated such that all instructions preceding the instruction causing the exception appear to have completed with respect to the executing thread.
3. The instruction causing the exception may appear not to have begun execution (except for causing the exception), may have been partially executed, or may have completed, depending on the interrupt type.
4. Architecturally, no subsequent instruction has begun execution.

### 6.4.2 Imprecise Interrupt

This architecture defines one imprecise interrupt, the Imprecise Mode Floating-Point Enabled Exception type Program interrupt.

When an Imprecise Mode Floating-Point Enabled Exception type Program interrupt occurs, the following conditions exist at the interrupt point.

1. SRRO addresses either the instruction causing the exception or some instruction following that instruction; see Section 6.5.9, "Program Interrupt" on page 1071.
2. An interrupt is generated such that all instructions preceding the instruction addressed by SRR0 appear to have completed with respect to the executing thread.
3. The instruction addressed by SRR0 may appear not to have begun execution (except, in some cases, for causing the interrupt to occur), may

## Version 3.0

have been partially executed, or may have completed; see Section 6.5.9.
4. No instruction following the instruction addressed by SRR0 appears to have begun execution, except that if an mtspr sequence started by mtgsr is active when the interrupt occurs, some of the sequence's mtsprs beyond the interrupt point may have been executed; see Chapter 11.

All Floating-Point Enabled Exception type Program interrupts are maskable using the MSR bits FEO and FE1. Although these interrupts are maskable, they differ significantly from the other maskable interrupts in that the masking of these interrupts is usually controlled by the application program, whereas the masking of all other maskable interrupts is controlled by either the operating system or the hypervisor.

### 6.4.3 Interrupt Processing

Associated with each kind of interrupt is an interrupt vector, which contains the initial sequence of instructions that is executed when the corresponding interrupt occurs.

Interrupt processing consists of saving a small part of the thread's state in certain registers, identifying the cause of the interrupt in other registers, and continuing execution at the corresponding interrupt vector location. When an exception exists that will cause an interrupt to be generated and it has been determined that the interrupt will occur, the following actions are performed. The handling of Machine Check interrupts (see Section 6.5.2) and System Call Vectored interrupts (see Section 6.5.27) differs from the description given below in several respects.

1. SRRO or HSRRO is loaded with an instruction address that depends on the type of interrupt; see the specific interrupt description for details.
2. Bits $33: 36$ and $42: 47$ of SRR1 or HSRR1 are loaded with information specific to the interrupt type.
3. Bits $0: 32,37: 41$, and $48: 63$ of SRR1 or HSRR1 are loaded with a copy of the corresponding bits of the MSR.
4. The MSR is set as shown in Figure 64 on page 1061. In particular, MSR bits IR and DR are set as specified by LPCR AIL (see Section 2.2), and MSR bit SF is set to 1 , selecting 64-bit mode. The new values take effect beginning with the first instruction executed following the interrupt.
5. Instruction fetch and execution resumes, using the new MSR value, at the effective address specific to the interrupt type. These effective addresses are shown in Figure 65 on page 1062. An offset may be applied to get the effective addresses, as specified by LPCR $_{\text {AIL }}$ (see Section 2.2).
Interrupts do not clear reservations obtained with Ibarx, Iharx, Iwarx, Idarx, or Iqarx.

## Programming Note

For instruction-caused interrupts, in some cases it may be desirable for the operating system to emulate the instruction that caused the interrupt, while in other cases it may be desirable for the operating system not to emulate the instruction. The following list, while not complete, illustrates criteria by which decisions regarding emulation should be made. The list applies to general execution environments; it does not necessarily apply to special environments such as program debugging, bring-up, etc.

In general, the instruction should be emulated if:

- The interrupt is caused by a condition for which the instruction description (including related material such as the introduction to the section describing the instruction) implies that the instruction works correctly. Example: Alignment interrupt caused by Imw for which the storage operand is not aligned, or by dcbz for which the storage operand is in storage that is Write Through Required or Caching Inhibited.
- The instruction is an illegal instruction that should appear, to the program executing it, as if it were supported by the implementation. Example: A Hypervisor Emulation Assistance interrupt is caused by an instruction that has been phased out of the architecture but is still used by some programs that the operating system supports.

If the instruction is a Storage Access instruction, the emulation must satisfy the atomicity requirements described in Section 1.4 of Book II.

In general, the instruction should not be emulated if:

- The purpose of the instruction is to cause an interrupt. Example: System Call interrupt caused by sc.
- The interrupt is caused by a condition that is stated, in the instruction description, potentially to cause the interrupt. Example: Alignment interrupt caused by Iwarx for which the storage operand is not aligned.
- The program is attempting to perform a function that it should not be permitted to perform. Example: Data Storage interrupt caused by Iwz for which the storage operand is in storage that the program should not be permitted to access. (If the function is one that the program should be permitted to perform, the conditions that caused the interrupt should be corrected and the program re-dispatched such that the instruction will be re-executed. Example: Data Storage interrupt caused by Iwz for which the storage operand is in storage that the program should be permitted to access but for which there currently is no PTE that satisfies the Page Table search.)


#### Abstract

Programming Note If a program modifies an instruction that it or another program will subsequently execute and the execution of the instruction causes an interrupt, the state of storage and the content of some registers may appear to be inconsistent to the interrupt handler program. For example, this could be the result of one program executing an instruction that causes a Hypervisor Emulation Assistance interrupt just before another instance of the same program stores an Add Immediate instruction in that storage location. To the interrupt handler code, it would appear that a hardware generated the interrupt as the result of executing a valid instruction.


## Programming Note

In order to handle Machine Check and System Reset interrupts correctly, the operating system should manage $\mathrm{MSR}_{\mathrm{RI}}$ as follows.

- In the Machine Check and System Reset interrupt handlers, interpret SRR1 bit 62 (where $\mathrm{MSR}_{\mathrm{RI}}$ is placed) as:
- 0 : interrupt is not recoverable
- 1 : interrupt is recoverable
- In each interrupt handler, when enough state has been saved that a Machine Check or System Reset interrupt can be recovered from, set $\mathrm{MSR}_{\mathrm{RI}}$ to 1 .
- In each interrupt handler, do the following (in order) just before returning.

1. Set $M_{\text {RRI }}$ to 0 .
2. Set SRR0 and SRR1 to the values to be used by rfid. The new value of SRR1 should have bit 62 set to 1 (which will happen naturally if SRR1 is restored to the value saved there by the interrupt, because the interrupt handler will not be executing this sequence unless the interrupt is recoverable).
3. Execute rfid.

For interrupts that set the SRRs other than Machine Check or System Reset, $\mathrm{MSR}_{\text {RI }}$ can be managed similarly when these interrupts occur within interrupt handlers for other interrupts that set the SRRs.

This Note does not apply to interrupts that set the HSRRs because these interrupts put the thread into hypervisor state, and either do not occur or can be prevented from occurring within interrupt handlers for other interrupts that set the HSRRs.

### 6.4.4 Implicit alteration of HSRR0 and HSRR1

Executing some of the more complex instructions may have the side effect of altering the contents of HSRR0 and HSRR1. The instructions listed below are guaranteed not to have this side effect. Any omission of instruction suffixes is significant; e.g., add is listed but add. is excluded.

## 1. Branch instructions

$b[\Omega[a], b c[\Lambda[a], b c I r[I], b c c t r[/]$
2. Fixed-Point Load and Store Instructions
lbz, Ibzx, Ihz, Ihzx, Iwz, Iwzx, Id, Idx, stb, stbx, sth, sthx, stw, stwx, std, stdx

Execution of these instructions is guaranteed not to have the side effect of altering HSRRO and HSRR1 only if the storage operand is aligned and $M_{\text {SR }}=0$.
3. Arithmetic instructions
addi, addis, add, subf, neg
4. Compare instructions
cmpi, cmp, cmpli, cmpl
5. Logical and Extend Sign instructions
ori, oris, xori, xoris, and, or, xor, nand, nor, eqv, andc, orc, extsb, extsh, extsw
6. Rotate and Shift instructions
rldicl, rldicr, rldic, rlwinm, rldcl, rldcr, rlwnm, rldimi, rlwimi, sld, slw, srd, srw
7. Other instructions

## isync

rfid, hrfid
mtspr, mfspr, mtmsrd, mfmsr

## Programming Note

Instructions excluded from the list include the following.

■ instructions that set or use XER CA

- instructions that set XER ${ }_{\mathrm{OV}}$ or XER $_{\mathrm{SO}}$
- andi., andis., and fixed-point instructions with Rc=1 (Fixed-point instructions with Rc=1 can be replaced by the corresponding instruction with Rc=0 followed by a Compare instruction.)
- all floating-point instructions
- mftb

These instructions, and the other excluded instructions, may be implemented with the assistance of the Hypervisor Emulation Assistance interrupt, or of implementation-specific interrupts that modify HSRR0 and HSRR1. The included instructions are guaranteed not to be implemented thus. (The included instructions are sufficiently simple as to be unlikely to need such assistance. Moreover, they are likely to be needed in interrupt handlers before HSRR0 and HSRR1 have been saved or after HSRR0 and HSRR1 have been restored.)

Similarly, fetching instructions may have the side effect of altering the contents of HSRR0 and HSRR1 unless $M S R_{I_{R}}=0$.

## Version 3.0

### 6.5 Interrupt Definitions

Figure 64 shows all the types of interrupts and the values assigned to the MSR for each. Figure 65 shows the effective address of the interrupt vector for each interrupt type. (Section 5.7.5 on page 987 summarizes all architecturally defined uses of effective addresses, including those implied by Figure 65.)

| Interrupt Type | MSR Bit |  |  |  |  |  |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| IR DR FE0 |  |  |  |  |  |  |  |  |
| SE1 | EE RI ME HV |  |  |  |  |  |  |  |$|$


| Interrupt Type | MSR Bit IR DR FE0 FE1 EE RI ME HV |
| :---: | :---: |
| 0 bit is set to 0 |  |
| 1 bit is set |  |
| - bit is not altered |  |
| $r$ for interrupts that are taken as if $\operatorname{LPCR}_{\text {AIL }}=3$, and for interrupts for which $\operatorname{LPCR}_{\text {AIL }}$ applies, if LPCR |  |
| p if the inte | wise not altered |
| e if LPES $=0$, set to 1 ; otherwise not altered |  |
| h if LPES $=1$, set to 0 ; otherwise not altered |  |
| $s \quad$ if LEV $=1$, set to 1; otherwise not altered |  |
| Settings for Other Bits |  |
| Bits BE, FP, PR, SE, TM, VEC, VSX, PMM, and bit 5 are set to 0 . |  |
| TM, FP, SLE, VEC, and VSX are set to 0 . |  |
| If the interrupt results in HV being equal to 1, the LE bit is copied from the HILE bit; otherwise the LE bit is copied from the LPCR ${ }_{\text {ILE }}$ bit. |  |
| The SF bit is set to 1 . |  |
| If the TS field contained Ob10 (Transactional) when the interrupt occurred, the TS field is se to 0b01 (Suspended); otherwise the TS field is not altered. |  |
| Reserved bits are set as if written as 0 . |  |

Figure 64. MSR setting due to interrupt

| Effective Address ${ }^{1}$ | Interrupt Type |
| :---: | :---: |
| 00. . 0000_0100 | System Reset |
| 00. . 0000_0200 | Machine Check |
| 00. . 0000_0300 | Data Storage |
| 00..0000_0380 | Data Segment |
| 00..0000_0400 | Instruction Storage |
| 00. . 0000_0480 | Instruction Segment |
| 00..0000_0500 | External |
| 00. .0000_0600 | Alignment |
| 00. . 0000_0700 | Program |
| 00. . 0000_0800 | Floating-Point Unavailable |
| 00. .0000_0900 | Decrementer |
| 00. . 0000_0980 | Hypervisor Decrementer |
| 00. . 0000_0A00 | Directed Privileged Doorbell |
| 00. . 0000_0B00 | Reserved |
| 00. . 0000_0C00 | System Call |
| 00. . 0000_0D00 | Trace |
| 00. . 0000_0E00 | Hypervisor Data Storage |
| 00. . 0000_0E20 | Hypervisor Instruction Storage |
| 00. . 0000_0E40 | Hypervisor Emulation Assistance |
| 00. . 0000_0E60 | Hypervisor Maintenance |
| 00. . 0000_0E80 | Directed Hypervisor Doorbell |
| 00. . 0000_0EA0 | Hypervisor Virtualization |
| 00. . 0000_0EC0 | Reserved |
| 00..0000_0EE0 | Reserved for implementa-tion-dependent interrupt for performance monitoring |
| 00..0000_0F00 | Performance Monitor |
| 00..0000_0F20 | Vector Unavailable |
| 00..0000_0F40 | VSX Unavailable |
| 00. . 0000_0F60 | Facility Unavailable |
| 00. . 0000_0F80 | Hypervisor Facility Unavailable |
| 00..0000_0FA0 | Reserved |
| 00. .0000_0FFF | Reserved |
| 00. . 0001_7000 | System Call Vectored |
| 00. . 0001_7020 | System Call Vectored |
| 00. . 0001_7FE0 | System Call Vectored |
| 00..0001_7FFF | (end of scv interrupt vectors) |



1 The values in the Effective Address column are interpreted as follows.

- 00...0000_0nnn means
$0 \times 0000 \_0000 \_0000 \_$Onnn unless the values of LPCR $_{\text {AIL }}$ and MSR $_{\text {HV IR DR }}$ cause the application of an effective address offset. See the description of LPCR ${ }_{\text {AIL }}$ in Section 2.2 for more details.
■ 0...00_0001_7nnn means 0x0000_0000_0001_7nnn unless the values of LPCR ${ }_{\text {AIL }}$ and MSR $_{\text {HV IR DR }}$ cause the usage of an alternate effective address. See the description of LPCR $_{\text {AIL }}$ in Section 2.2 for details.
2 Effective addresses 0x0000_0000_0000_0000 through 0x0000_0000_0000_00FF are used by software and will not be assigned as interrupt vectors.

Figure 65. Effective address of interrupt vector by interrupt type

## Programming Note

When address translation is disabled, use of any of the effective addresses that are shown as reserved in Figure 65 risks incompatibility with future implementations.

### 6.5.1 System Reset Interrupt

If a System Reset exception causes an interrupt that is not context synchronizing or causes the loss of a Machine Check exception or a Direct External exception, or if the state of the thread has been corrupted, the interrupt is not recoverable.
When the thread is in any power-saving level, a System Reset interrupt occurs when a System Reset exception exists. When the thread is in a power-saving level that was entered when PSSCR ${ }_{E C}=1$, a System Reset interrupt also occurs when any of the following events occurs provided that the event is enabled to cause exit from power-saving mode (see Section 2.2). When the thread is in a power-saving level that allows the state of the LPCR to be lost, it is implementation-specific whether the following events, when enabled, cause exit, or whether only a system-reset exception causes exit.

- External
- Decrementer
- Directed Privileged Doorbell
- Directed Hypervisor Doorbell
- Hypervisor Maintenance
- ■ Hypervisor Virtualization exception
- Implementation-specific

SRR1 indicates the exception that caused exit from power-saving mode as specified below.

The following registers are set:
SRRO If the interrupt did not occur when the thread was in power-saving mode, set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present; otherwise, set to an undefined value.
If the interrupt occurred while the thread was in power-saving mode, set to the effective address of the instruction following the stop instruction when stop is executed with PSSCR bit ESL=0 and fields RL, MTL, and PSLL set to values that do not allow state loss; otherwise, set to an undefined value.

## Programming Note

Whenever stop is executed in privileged non-hypervisor state, the hypervisor typically sets both PSSCR $_{\text {ESL }}$ and PSSCR $_{\text {EC }}$ to 0 , and sets RL and MTL to values that do not cause state loss. If an interrupt causes exit to power-saving mode (either because the interrupt was a System Reset or Machine Check interrupt or $\mathrm{MSR}_{E E}=1$ ), then SRRO for that interrupt contains the effective address of the instruction immediately following stop.

## SRR1

33 Implementation-dependent.
34:36 Set to 0 .
42:45 If the interrupt did not occur when the thread was in power-saving mode, set to an implementation-specific value. If the interrupt occurred when the thread was in power-saving mode, set to indicate the
exception that caused exit from power-saving mode as shown below:

| SRR1 $_{42: 45}$ | Exception |
| :--- | :--- |
| 0000 | Reserved |
| 0001 | Reserved |
| 0010 | Implementation specific |
| 0011 | Directed Hypvsr Doorbell |
| 0100 | System Reset |
| 0101 | Directed Privlgd Doorbell |
| 0110 | Decrementer |
| 0111 | Reserved |
| 1000 | External |
| 1001 | Hypervisor Virtualization |
| 1010 | Hypervisor Maintenance |
| 1011 | Reserved |
| 1100 | Implementation specific |
| 1101 | Reserved |
| 1110 | Implementation specific |
| 1111 | Reserved |

If multiple events that cause exit from power-saving mode exist, the event reported is the exception corresponding to the interrupt that would have occurred if the same conditions existed and the thread was not in power-saving mode.

46:47 Set to indicate whether the interrupt occurred when the thread was in power-saving mode and, if so, the extent to which resource state was maintained while the thread was in power-saving mode, as follows:

00 The interrupt did not occur when the thread was in power-saving mode.

01 The interrupt occurred when the thread was in power-saving mode. The state of all resources was maintained as if the thread was not in power-saving mode.

10 The interrupt occurred when the thread was in power-saving mode. The state of some resources was not maintained, but the state of all hypervisor resources, including the TB, PURR, and SPURR, was maintained as if the thread was not in power-saving mode and the state of all other resources is such that the hypervisor can resume execution. (See Section 2.6 for the list of hypervisor resources.)

11 The interrupt occurred when the thread was in power-saving mode. The state of some resources was not maintained, and the state of some hypervisor resources was not maintained or the state of some resources is such that the hypervisor cannot resume execution

## Programming Note

Although the resources that are maintained in power-saving levels that allow loss of state are implementa-tion-dependent, the hypervisor can avoid implementation-dependence in the portion of the System Reset and Machine Check interrupt handlers that recover from having been in power-saving mode by using the contents of SRR1 ${ }_{46: 47}$, to determine what state to restore. (To avoid implementa-tion-dependence, the hypervisor must assume that only the resources indicated in SRR1 ${ }_{46: 47}$ have been preserved.

62 If the interrupt did not occur while the thread was in a power-saving level that was entered when PSSCR $_{E C}=1$, loaded from bit 62 of the MSR if the thread is in a recoverable state; otherwise set to 0 . If the interrupt occurred while the thread was in a power-saving level that was entered when PSSCR $_{E C}=1$, set to 1 if the thread is in a recoverable state; otherwise set to 0 .
Others Loaded from the MSR.

## MSR See Figure 64 on page 1061.

In addition, if the interrupt occurs when the thread is in a power-saving level that was entered when PSS$\mathrm{CR}_{\mathrm{EC}}=1$ and is caused by an exception other than a System Reset exception, all other registers, except HSRRO and HSRR1, that would be set by the corresponding interrupt if the exception occurred when the thread was not in power-saving mode are set by the System Reset interrupt, and are set to the values to which they would be set if the exception occurred when the thread was not in power-saving mode.

Execution resumes at effective address 0x0000_0000_0000_0100.

The means for software to distinguish between power-on Reset and other types of System Reset are implementation-dependent.

### 6.5.2 Machine Check Interrupt

The causes of Machine Check interrupts are implemen-tation-dependent. For example, a Machine Check interrupt may be caused by a reference to a storage location that contains an uncorrectable error or does not exist (see Section 5.6), or by an error in the storage subsystem.

When the thread is not in power-saving mode, Machine Check interrupts are enabled when $\mathrm{MSR}_{\mathrm{ME}}=1$; if $M S R_{M E}=0$ and a Machine Check exception occurs, the thread enters the Checkstop state. When the thread is in a power-saving level that does not allow loss of hypervisor state, Machine Check interrupts are treated as enabled when $\mathrm{LPCR}_{51}=1$ and cannot occur when $\operatorname{LPCR}_{51}=0$. When the thread is in a power-saving level that allows loss of hypervisor state, it is implementa-tion-specific whether Machine Check interrupts are | treated as enabled $\operatorname{LPCR}_{51}=1$ or if they cannot occur. If a Machine Check exception occurs while the thread is in power-saving mode and the Machine Check exception is not enabled to cause exit from power-saving mode, the result is implementation specific.

The Checkstop state may also be entered if an access is attempted to a storage location that does not exist (see Section 5.6), or if an implementation-dependent hardware error occurs that prevents continued operation.

## Disabled Machine Check (Checkstop State)

When a thread is in Checkstop state, instruction processing is suspended and generally cannot be restarted without resetting the thread. Some implementations may preserve some or all of the internal state of the thread when entering Checkstop state, so that the state can be analyzed as an aid in problem determination.

## Enabled Machine Check

If a Machine Check exception causes an interrupt that is not context synchronizing or causes the loss of a Direct External exception, or if the state of the thread has been corrupted, the interrupt is not recoverable.

The following registers are set:
SRRO If the interrupt occurred when the thread was in a power-saving mode that was entered with PSSCR bit ESL=0, and fields RL, MTL, and PSLL set to values that do not allow state loss; set on a "best effort" basis to the effective address of some instruction that was executing or was about to be executed when the Machine Check exception occurred; otherwise set to an undefined value.

46:47 Set to indicate whether the interrupt occurred when the thread was in power-saving mode and, if so, the extent to which resource state was maintained while the thread was in power-saving mode, as follows.

00 The interrupt did not occur when the thread was in power-saving mode.

01 The interrupt occurred when the thread was in power-saving mode. The state of all resources was maintained as if the thread was not in power-saving mode.

The interrupt occurred when the thread was in power-saving mode. The state of some resources was not maintained, but the state of all hypervisor resources, including the TB, PURR, and SPURR, was maintained as if the thread was not in power-saving mode and the state of all other resources is such that the hypervisor can resume execution. (See Section 2.6 for the list of hypervisor resources.)

11 The interrupt occurred when the thread was in power-saving mode. The state of some resources was not maintained, and the state of some hypervisor resources was not maintained or the state of some resources is such that the hypervisor cannot resume execution.

## Programming Note

Although the resources that are maintained in power-saving mode (except when all resources are maintained) are implementation-dependent, the hypervisor can avoid implementa-tion-dependence in the portion of the System Reset and Machine Check interrupt handlers that recover from having been in power-saving mode by using the contents of SRR1 ${ }_{46: 47}$, to determine what state to restore. (To avoid implementation-dependence in the portion of the hypervisor that enters power-saving mode, the hypervisor must use the specification of the four instructions to determine what state to save.)

62 If the interrupt did not occur while the thread was in a power-saving level that was entered when PSSCR ${ }_{E C}=1$, loaded from bit 62 of the MSR if the thread is in a recoverable state; otherwise set to 0 . If the interrupt occurred while the thread was in a power-saving level that was entered when $\operatorname{PSSCR}_{E C}=1$, set to 1 if the thread is in a recoverable state; otherwise set to 0 .

Others Set to an implementation-dependent value.
MSR See Figure 64.
DSISR Set to an implementation-dependent value.
DAR Set to an implementation-dependent value.
I ASDR Set to an implementation-dependent value.
Execution resumes at effective address 0x0000_0000_0000_0200.

A Machine Check interrupt caused by the existence of multiple SLB entries or TLB entries (or similar entries in implementation-specific translation caches) which translate a given effective or virtual address (see Sections 5.7.8.2 and 5.7.9.2.) must occur while still in the context of the partition that caused it. The interrupt must be presented in a way that permits continuing execution, with damage limited to the causing partition. Treating the exception as instruction-caused will achieve these requirements.

## Programming Note

If a Machine Check interrupt is caused by an error in the storage subsystem, the storage subsystem may return incorrect data, which may be placed into registers. This corruption of register contents may occur even if the interrupt is recoverable.

### 6.5.3 Data Storage Interrupt

A Data Storage interrupt occurs when no higher priority exception exists and either
(a) HRIIGR=0b00, the value of the expression

$$
\left(\left(\mathrm{MSR}_{\mathrm{HV} \mathrm{PR}}=0 \mathrm{~b} 10\right) \mid\left(\left(\neg \mathrm{VPM}_{1} \mid \neg \mathrm{PRTE}_{\mathrm{V}}\right) \& \mathrm{MSR}_{\mathrm{DR}}\right)\right)
$$

is 1 , and a data access cannot be performed,
except for the case of $M_{\text {M }} R_{H V} \neq 0 b 10$,
$V P M_{1}=0, L P C R_{K B V}=1$, and a Virtual Storage
Page Class Key Protection exception exists or
(b) HRIIGR $\neq 0$ b00 and process-scoped translation
prevents a data access from being performed
for any of the following reasons. (In the expression for (a) above, " $\neg \mathrm{PRTE}_{V} "$ is shorthand representing the case of an invalid segment table descriptor stopping the translation process.)

- Data address translation is enabled $\left(\mathrm{MSR}_{\mathrm{DR}}=1\right)$ and the virtual address of any byte of the storage location specified by a Load, Store, icbi, dcbz, dcbst, or dcbf[I] instruction cannot be translated to a real address because no valid PTE was found for the process-scoped translation or paravirtualized translation with VPM off.
- The address of the appropriate process table entry or segment table entry group cannot be translated when HRIIGR=0b00 and either $\mathrm{VPM}_{1}=0$ or the process table entry is invalid (independent of VPM ${ }_{1}$ ).
- The effective address specified by a Iq, stq, Iwat, Idat, Ibarx, Iharx, Iwarx, Idarx, Iqarx, stwat, stdat, stbcx., sthcx., stwcx., stdcx., or stqcx. instruction refers to storage that is Write Through Required or Caching Inhibited; or the effective address specified by a Iwat, Idat, stwat, or stdat instruction refers to storage that is Guarded.
- The effective address specified by a copy or paste[.] instruction refers to storage that is Caching Inhibited.
- The effective address specified by an instruction other than copy or paste[.] refers to CSM.
- An accelerator is specified as the source of a copy instruction or an attempt is made to access an accelerator that is not properly initialized for the software's use.
- A transfer specified to access CSM does not specify (local) main storage as the other end of the transfer.
- The access violates Basic Storage Protection.
- The access violates Virtual Page Class Key Storage Protection and $\mathrm{LPCR}_{\mathrm{KBV}}=0$.
- The process- and partition-scoped page attributes conflict.
- An unsupported radix tree configuration is found in the process-scoped tables.
- A reference or change bit update cannot be performed in a process-scoped PTE.
- A Data Address Watchpoint match occurs.
- An attempt is made to execute a Fixed-Point Load or Store Caching Inhibited instruction with $\mathrm{MSR}_{\mathrm{DR}}=1$ or specifying a storage location that is specified by the Hypervisor Real Mode Storage Control facility to be treated as non-Guarded.

A Data Storage interrupt also occurs when no higher priority exception exists and an attempt is made to execute a Load Atomic or Store Atomic instruction specifying an invalid function code.

If a stbcx., sthcx., stwcx., stdcx., or stqcx. would not perform its store in the absence of a Data Storage interrupt, and either (a) the specified effective address refers to storage that is Write Through Required or Caching Inhibited, or (b) a non-conditional Store to the specified effective address would cause a Data Storage interrupt, it is implementation-dependent whether a Data Storage interrupt occurs.

If the XER specifies a length of zero for an indexed Move Assist instruction, a Data Storage interrupt does not occur.

The following registers are set:
SRRO Set to the effective address of the instruction that caused the interrupt.
SRR1
33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64.
DSISR
32
33 Set to 1 if $M S R_{D R}=1$ and the translation for an attempted access is not found in the Page Table; otherwise set to 0 ..
34 Set to 1 if the process- and parti-tion-scoped page attributes conflict; otherwise set to 0 .
Set to 0 .
35
36 Set to 1 if the access is not permitted by Figure 43 45, or the privilege, read, or read/write bits in Figure 44 as appropriate; otherwise set to 0 .
37 Set to 1 if the access is due to a Iq, stq, Iwat, Idat, Ibarx, Iharx, Iwarx, Idarx, Iqarx, stwat, stdat, stbcx., sthcx., stwcx., stdcx., or stqcx. instruction that addresses storage that is Write Through Required or Caching Inhibited; or if the access is due to a Iwat, Idat, stwat, or stdat instruction that addresses storage that is Guarded; or if the access is due to a copy or paste[.] instruction that addreses storage that is caching inhibited; or if the access is due to an instruction other than copy or paste[.] addressing CSM; otherwise set to 0 .
Set to 1 for a Store or dcbz instruction; otherwise set to 0 .
39:40 Set to 0 .
41 Set to 1 if a Data Address Watchpoint match occurs; otherwise set to 0 .
42 Set to 1 if the access is not permitted by virtual page class key protection; otherwise set to 0 .

Set to 0 .
Set to 1 if an unsupported radix tree configuration is found during the translation process; otherwise set to 0 .
45 Set to 1 if an attempt to atomically set a reference or change bit fails; otherwise set to 0.

46 Set to 1 if the address of the appropriate process table entry or segment table entry group cannot be translated when VPM $_{1}=0$ and $H R \| G R=0 b 00$, or the process table
entry is invalid (independent of $\mathrm{VPM}_{1}$ ) when HR\|GR=0b00.
47:58 Set to 0 .
59 Set to 1 if a transfer specified to access CSM does not specify (local) main storage as the other end of the transfer; otherwise set to 0 .
60 Set to 1 if an accelerator is specified as the source of a copy instruction or an attempt is made to access an accelerator that is not properly initialized for the software's use; otherwise set to 0 .

## Programming Note

The exceptions identified by bits 59 and 60 are regarded as fatal programming errors. Additional information may be retained by the platform for the second cause of bit 60 .

61 Set to 1 if an attempt is made to execute a Load Atomic or Store Atomic instruction specifying an invalid function code; otherwise set to 0 .
62 Set to 1 if an attempt is made to execute a Fixed-Point Load or Store Caching Inhibited instruction with $M_{\text {DR }}=1$ or specifying a storage location that is specified by the Hypervisor Real Mode Storage Control facility to be treated as non-Guarded.
63 Set to 0 .
DAR Set to the effective address of a storage element as described in the following list. The list should be read from the top down; the DAR is set as described by the first item that corresponds to an exception that is reported in the DSISR. For example, if a Load Word instruction causes a storage protection violation and a Data Address Watchpoint match (and both are reported in the DSISR), the DAR is set to the effective address of a byte in the first aligned doubleword for which access was attempted in the page that caused the exception.

- a Data Storage exception occurs for reasons other than a Data Address Watchpoint match
- a byte in the block that caused the exception, for a Cache Management instruction
- a byte in the first aligned quadword for which access was attempted in the page that caused the exception, for a quadword Load or Store instruction (i.e., a Load or Store instruction for which the storage operand is a quadword; "first" refers to address order: see Section 6.7)
- a byte in the first aligned doubleword for which access was attempted in the page that caused the exception, for a non-quadword Load or Store instruction
- undefined, for a Data Address Watchpoint match

For the cases in which the DAR is specified above to be set to a defined value, if the interrupt occurs in 32-bit mode the high-order 32 bits of the DAR are set to 0 .

If multiple Data Storage exceptions occur for a given effective address, any one or more of the bits corresponding to these exceptions may be set to 1 in the DSISR. However, if one or more DSI-causing exceptions occur together with a Virtualized Page Class Key Storage Protection exception that occurs when $\mathrm{LPCR}_{\text {KBV }}=1$ and Virtualized Partition Memory is disabled by $V^{2} M_{1}=0$, an HDSI results, and all of the exceptions are reported in the HDSISR.
Execution resumes at effective address 0x0000_0000_0000_0300, possibly offset as specified in Figure 65.

### 6.5.4 Data Segment Interrupt

| For Paravirtualized HPT Translation, a Data Segment interrupt occurs when no higher priority exception exists and a data access cannot be performed because data address translation is enabled and the effective address of any byte of the storage location specified by a Load, Store, icbi, dcbz, dcbst, or dcbf[] instruction cannot be translated to a virtual address.
For Radix Tree Translation (in other than hypervisor real mode), a Data Segment interrupt occurs when no higher priority exception exists and a data access cannot be performed because for the effective address specified by a Load, Store, icbi, dcbz, dcbst, or dcbf[]] instruction, $\mathrm{EA}_{0: 1}=0 \mathrm{~b} 01$ or $\mathrm{EA}_{0: 1}=0 \mathrm{~b} 10$ when $\mathrm{MSR}_{\mathrm{HV}}$ $\mathrm{PR} \neq 0 \mathrm{~b} 10$ and data address translation is enabled, or $\mathrm{EA}_{2: 63}$ is outside the range translated by the appropriate Radix Tree.

If a stbcx., sthcx., stwcx., stdcx., or stqcx. would not perform its store in the absence of a Data Segment interrupt and a non-conditional Store to the specified effective address would cause a Data Segment interrupt, it is implementation-dependent whether a Data Segment interrupt occurs.
If the XER specifies a length of zero for an indexed Move Assist instruction, a Data Segment interrupt does not occur.
The following registers are set:
SRRO Set to the effective address of the instruction that caused the interrupt.
SRR1

33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.

## MSR See Figure 64.

DSISR Set to an undefined value.
DAR Set to the effective address of a storage element as described in the following list.

- a byte in the block that caused the exception, for a Cache Management instruction
- a byte in the first aligned quadword for which access was attempted in the segment that caused the exception, for a quadword Load or Store instruction (i.e., a Load or Store instruction for which the storage operand is a quadword; "first" refers to address order: see Section 6.7)
- a byte in the first aligned doubleword for which access was attempted in the segment that caused the exception, for a non-quadword Load or Store instruction
If the interrupt occurs in 32-bit mode the high-order 32 bits of the DAR are set to 0 .
Execution resumes at effective address 0x0000_0000_0000_0380, possibly offset as specified in Figure 65.


## Programming Note

A Data Segment interrupt occurs if $\mathrm{MSR}_{\mathrm{DR}}=1$ and the translation of the effective address of any byte of the specified storage location is not found in the SLB (or in any implementation-specific address translation lookaside information).

### 6.5.5 Instruction Storage Interrupt

An Instruction Storage interrupt occurs when no higher priority exception exists and either
(a) $\mathrm{HR} \| \mathrm{GR}=0 \mathrm{~b} 00$, the value of the expression $\left(\left(\mathrm{MSR}_{\mathrm{HV} \text { PR }}=0 \mathrm{~b} 10\right) \mid\left(\left(\neg \mathrm{VPM}_{1} \mid \neg \mathrm{PRTE}_{\mathrm{V}}\right) \& \mathrm{MSR}_{\mathrm{IR}}\right)\right)$ is 1 , and the next instruction to be executed cannot be fetched, or
(b) $\mathrm{HR} \| G R \neq 0 \mathrm{O} 00$ and process-scoped translation prevents the next instruction to be executed from being fetched
for any of the following reasons. (In the expression for (a) above, " $\neg P R T E_{V}$ " is shorthand representing the case of an invalid segment table descriptor stopping the translation process.)

- Instruction address translation is enabled and the virtual address cannot be translated to a real address because no valid PTE was found for the process-scoped translation or paravirtualized translation with VPM off.
- The address of the appropriate process table entry or segment table entry group cannot be translated when $\mathrm{HR} \| \mathrm{GR}=0 \mathrm{~b} 00$ and either $\mathrm{VPM}_{1}=0$ or the process table entry is invalid (independent of VPM ${ }_{1}$ ).
- The fetch access violates storage protection.
- The process- and partition-scoped page attributes conflict.
- An unsupported radix tree configuration is found in the process-scoped tables.
- A reference bit update cannot be performed in a process-scoped PTE.
The following registers are set:
SRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present (if the interrupt occurs on attempting to fetch a branch target, SRRO is set to the branch target address).


## SRR1

33
Set to 1 if $\mathrm{MSR}_{\mathrm{IR}^{\prime}}=1$ and the translation for an attempted access is not found in the Page Table; otherwise set to 0 .
34 Set to 1 if the process- and parti-tion-scoped page attributes conflict; otherwise set to 0 .
35 Set to 1 if the access is to No-execute (as indicated by the N bit in the segment table entry or the $N$ bit in the HPT PTE or the Execute and Privilege bits in the EAA field of the Radix PTE and IAMR key 0 ) or Guarded storage; otherwise set to 0 .
36 Set to 1 if the access is not permitted by Figure 43 or 45 , as appropriate; otherwise set to 0 .

42 Set to 1 if the access is not permitted by virtual page class key protection; otherwise set to 0 .
43 Set to 0.
44 Set to 1 if an unsupported radix tree configuration is found during the translation process; otherwise set to 0 .
45 Set to 1 if an attempt to atomically set a reference bit fails; otherwise set to 0 .
46 Set to 1 if the address of the appropriate process table entry or segment table entry group cannot be translated when VPM $_{1}=0$ and HR\|GR=0b00, or the process table entry is invalid (independent of $\mathrm{VPM}_{1}$ ) when HRIIGR=0b00.
47 Set to 0.
Others Loaded from the MSR.
MSR See Figure 64.

If multiple Instruction Storage exceptions occur due to attempting to fetch a single instruction, any one or more of the bits corresponding to these exceptions may be set to 1 in SRR1.

Execution resumes at effective address 0x0000_0000_0000_0400, possibly offset as specified in Figure 65.

### 6.5.6 Instruction Segment Interrupt

I For Paravirtualized HPT Translation, an Instruction Segment interrupt occurs when no higher priority exception exists and the next instruction to be executed cannot be fetched because instruction address translation is enabled and the effective address cannot be translated to a virtual address.

For Radix Tree Translation (in other than hypervisor real mode), an Instruction Segment interrupt occurs when no higher priority exception exists and the next instruction to be executed cannot be fetched because $E A_{0: 1}=0 \mathrm{~b} 01$ or $\mathrm{EA}_{0: 1}=0 \mathrm{~b} 10$ when $\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR} \neq 0 \mathrm{~b} 10$ and instruction address translation is enabled, or $\mathrm{EA}_{2: 63}$ is outside the range translated by the appropriate Radix Tree.

The following registers are set:
SRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present (if the interrupt occurs on attempting to fetch a branch target, SRRO is set to the branch target address).

## SRR1

33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0480, possibly offset as specified in Figure 65.

## Programming Note

An Instruction Segment interrupt occurs if $\mathrm{MSR}_{\mathrm{IR}}=1$ and the translation of the effective address of the next instruction to be executed is not found in the SLB (or in any implementation-specific address translation lookaside information).

### 6.5.7 External Interrupt

An External interrupt is classified as being either a Direct External interrupt or a Mediated External interrupt. Throughout this Book, usage of the phrase "External interrupt', without further classification, refers to
both a Direct External interrupt and a Mediated External interrupt.

### 6.5.7.1 Direct External Interrupt

A Direct External interrupt occurs when no higher priority exception exists, a Direct External exception exists, and the value of the expression
| $\mathrm{MSR}_{\text {EE }}$ \& $\neg\left(\mathrm{MSR}_{H V} \& \neg \mathrm{MSR}_{\text {PR }}\right.$ \& $\left.\mathrm{LPCR}_{\text {HEIC }}\right)$ I $\left(\neg(\mathrm{LPES}) \&\left(\neg\left(\mathrm{MSR}_{\mathrm{HV}}\right) \mid \mathrm{MSR}_{\mathrm{PR}}\right)\right)$
is one. The occurrence of the interrupt does not cause the exception to cease to exist.

## Programming Note

When HEIC=1, Direct External exceptions will not result in external interrupts when the processor is in hypervisor state. This enables the Hypervisor Interrupt Virtualization handler to prevent External interrupts from occurring during the Hypervisor Virtualization interrupt handler.

When LPES $=0$, the following registers are set:
HSRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present.

## HSRR1

33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
When LPES $=1$, the following registers are set:
SRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present.
SRR1
33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0500, possibly offset as specified in Figure 65.

## Programming Note

Because the value of $\mathrm{MSR}_{\text {EE }}$ is always 1 when the thread is in problem state, the simpler expression
$\mathrm{MSR}_{\text {EE }}$ \& $\neg\left(\mathrm{MSR}_{\mathrm{HV}}\right.$ \& $\neg \mathrm{MSR}_{\text {PR }}$ \& $\left.\mathrm{LPCR}_{\text {HEIC }}\right) ।$
$\neg\left(\right.$ LPES $\mid$ MSR $\left._{H V}\right)$
is equivalent to the expression given above.

## Programming Note

The Direct External exception has the same meaning as the External exception in versions of the architecture prior to Version 2.05.

### 6.5.7.2 Mediated External Interrupt

A Mediated External interrupt occurs when no higher priority exception exists, a Mediated External exception exists (see the definition of LPCR $_{\text {MER }}$ in Section 2.2), and the value of the expression

$$
\mathrm{MSR}_{E E} \&\left(\neg\left(\mathrm{MSR}_{H V}\right) \mid \mathrm{MSR}_{P R}\right)
$$

is one. The occurrence of the interrupt does not cause the exception to cease to exist.
When LPES $=0$, the following registers are set:
HSRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present.
HSRR1
33:36 Set to 0.
42 Set to 1.
43:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
When LPES=1, the following registers are set:
SRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present.
SRR1
33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0500, possibly offset as specified in Figure 65.

### 6.5.8 Alignment Interrupt

Many causes of Alignment interrupt involve storage operand alignment. Storage operand alignment is defined in Section 1.11.1 of Book I.

An Alignment interrupt occurs when no higher priority exception exists and an attempt is made to execute an instruction in a manner that is required, by the instruction description, to cause an Alignment interrupt. These cases are as follows.

- A Load/Store Multiple instruction that is executed in Little-Endian mode
- A Move Assist instruction that is executed in Lit-tle-Endian mode, unless the string length is zero
- A copy, paste[.], Iwat, Idat, Iharx, Iwarx, Idarx, Iqarx, stwat, stdat, sthcx., stwcx., stdcx., or stqcx. instruction that has an unaligned storage operand, unless execution of the instruction yields boundedly undefined results
- The operand(s) of a Load Atomic or Store Atomic instruction cross(es) a 32-byte boundary.
An Alignment interrupt may occur when no higher priority exception exists and a data access cannot be performed for any of the following reasons.
■ The storage operand of Ifdp, Ifdpx, stfdp, stfdpx, Ixsihzx, or stxsihx is unaligned.
- The storage operand of $\boldsymbol{I q}$ or $\boldsymbol{s t q}$ is unaligned.
- The storage operand of a Floating-Point Storage Access or VSX Storage Access instruction other than Ifdp, Ifdpx, stfdp, stfdpx, Ixsihzx, Ixsibzx, stxsihx, or stxsibx is not word-aligned.
- The storage operand of a Load/Store Multiple Word instruction is not word-aligned and the thread is in Big-Endian mode.
- The storage operand of a Load/Store Multiple Doubleword instruction is not doubleword-aligned and the thread is in Big-Endian mode.
- The storage operand of a Load/Store Multiple, Ifdp, Ifdpx, stfdp, stfdpx, or dcbz instruction is in storage that is Write Through Required or Caching Inhibited.
- The storage operand of a Move Assist instruction is in storage that is Write Through Required or Caching Inhibited and has length greater than zero.
- The storage operand of a Load or Store instruction is unaligned and is in storage that is Write Through Required or Caching Inhibited.
- The storage operand of a Storage Access instruction crosses a segment boundary, or crosses a boundary between virtual pages that have different storage control attributes.

The following registers are set:
SRRO Set to the effective address of the instruction that caused the interrupt.

## SRR1

33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64.

DAR Set to the effective address computed by the instruction, except that if the interrupt occurs in 32-bit mode the high-order 32 bits of the DAR are set to 0 .

Execution resumes at effective address 0x0000_0000_0000_0600, possibly offset as specified in Figure 65.

## Programming Note

If an Alignment interrupt occurs for a case in the second bulleted list above, the Alignment interrupt handler should emulate the instruction. The emulation must satisfy the atomicity requirements described in Section 1.4 of Book II.

If an Alignment interrupt occurs for a case in the first bulleted list above, the Alignment interrupt handler must not attempt to emulate the instruction, but instead should treat the instruction as a programming error.

### 6.5.9 Program Interrupt

A Program interrupt occurs when no higher priority exception exists and one of the following exceptions arises during execution of an instruction:

## Floating-Point Enabled Exception

A Floating-Point Enabled Exception type Program interrupt is generated when the value of the expression

$$
\left(\mathrm{MSR}_{\mathrm{FE} 0} \mid \mathrm{MSR}_{\mathrm{FE} 1}\right) \& \mathrm{FPSCR}_{\mathrm{FEX}}
$$

is 1 . FPSCR $_{\text {FEX }}$ is set to 1 by the execution of a floating-point instruction that causes an enabled exception, including the case of a Move To FPSCR instruction that causes an exception bit and the corresponding enable bit both to be 1 .

## TM Bad Thing

A TM Bad Thing type Program interrupt is generated when any of the following occurs.

- An rfebb, rfid, rfscv, hrfid, or mtmsrd instruction attempts to cause an illegal state transition (see Section 3.2.2).
- An rfid, rfscv, hrfid, or mtmsrd instruction attempts to cause a transition to Problem state with an active transaction (Transactional or Suspended state) when TM is disabled by the PCR ( $\mathrm{PCR}_{\mathrm{TM}}=1$ or $\mathrm{PCR}_{\mathrm{v} 2.06}=1$ ).
- An rfebb instruction in Problem state attempts to cause a transition to Transactional or Suspended state when $\mathrm{PCR}_{T M}=1$ (i.e., a latent non-zero TS value was in the BESCR).
- An attempt is made to execute trechkpt. in Transactional or Suspended state or when TEXASR $_{\text {FS }}=0$.
- An attempt is made to execute tend. in Suspended state.
- An attempt is made to execute treclaim. in Non-transactional state.
- An attempt is made to execute an mtspr instruction targeting a TM register in other
than Non-transactional state, with the exception of TFHAR in Suspended state.
- An attempt is made to execute a stop instruction in Suspended state.


## Privileged Instruction

The following applies if the instruction is executed when $M S R_{P R}=1$.

A Privileged Instruction type Program interrupt is generated when execution is attempted of a privileged instruction, or of an mtspr or mfspr instruction with an SPR field that contains a value having $\mathrm{spr}_{0}=1$.
The following applies if the instruction is executed when $M_{\text {HV }}$ PR $=0 b 00$ and LPCR EVIRT $=0$.

A Privileged Instruction type Program interrupt is generated when execution is attempted of an mtspr or mfspr instruction with an SPR field that designates an SPR that is accessible by the instruction only when the thread is in hypervisor state, or when execution of a hypervisor-privileged instruction is attempted.

## - Programming Note

These are the only cases in which a Privileged Instruction type Program interrupt can be generated when $M S R_{P R}=0$. They can be distinguished from other causes of Privileged Instruction type Program interrupts by examining $\mathrm{SRR}_{49}$ (the bit in which $M_{S R}{ }_{P R}$ was saved by the interrupt).

## Trap

A Trap type Program interrupt is generated when any of the conditions specified in a Trap instruction is met.

The following registers are set:
SRRO For all Program interrupts except a Float-ing-Point Enabled Exception type Program interrupt, set to the effective address of the instruction that caused the corresponding exception.
For a Floating-Point Enabled Exception type Program interrupt, set as described in the following list.

- If MSR ${ }_{\text {FE0 FE1 }}=0 b 00$, FPSCR $_{\text {FEX }}=1$, and an instruction is executed that changes $\mathrm{MSR}_{\text {FE0 FE1 }}$ to a nonzero value, set to the effective address of the instruction that the thread would have attempted
to execute next if no interrupt conditions were present.


## Programming Note

Recall that all instructions that can alter $\mathrm{MSR}_{\text {FE0 FE1 }}$ are context synchronizing, and therefore are not initiated until all preceding instructions have reported all exceptions they will cause.

- If $\mathrm{MSR}_{\text {FEO }} \mathrm{FE}=0 \mathrm{~b} 11$, set to the effective address of the instruction that caused the Floating-Point Enabled Exception.
- If $\mathrm{MSR}_{\text {FE0 }} \mathrm{FE}=0 \mathrm{~b} 01$ or $0 b 10$, set to the effective address of the first instruction that caused a Floating-Point Enabled Exception since the most recent time FPSCR $_{\text {FEX }}$ was changed from 1 to 0 or of some subsequent instruction.


## Programming Note

If SRRO is set to the effective address of a subsequent instruction, that instruction will not be beyond the first such instruction at which synchronization of floating-point instructions occurs. (Recall that such synchronization is caused by Floating-Point Status and Control Register instructions, as well as by execution synchronizing instructions and events.)

## SRR1

45 Set to 1 for a Privileged Instruction type Program interrupt; otherwise set to 0 .
46 Set to 1 for a Trap type Program interrupt; otherwise set to 0 .
47 Set to 0 if SRRO contains the address of the instruction causing the exception and there is only one such instruction; otherwise set to 1 .

## Programming Note

SRR1 ${ }_{47}$ can be set to 1 only if the exception is a Floating-Point Enabled Exception and either MSR FEO FE1 $=$ Ob01 or Ob10 or MSR FEO FE1 has just been changed from $0 b 00$ to a nonzero value. (SRR1 ${ }_{47}$ is always set to 1 in the last case.)

Others Loaded from the MSR.
Exactly one of bits $42,43,45$, and 46 is set to 1 .
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0700, possibly offset as specified in Figure 65.

## Programming Note

In versions of the architecture that precede V. 2.05, the conditions that now cause a Hypervisor Emulation Assistance interrupt instead caused an "lllegal Instruction type Program interrupt". This was a Program interrupt for which registers (SRRO, SRR1, and the MSR) were set as described above for the Privileged Instruction type Program interrupt, except that SRR1 ${ }_{44}$ was set to 1 and SRR1 ${ }_{45}$ was set to 0 . Thus operating systems have code to handle these conditions, at the Program interrupt vector location. For this reason, if a Hypervisor Emulation Assistance interrupt occurs, when the thread is not in hypervisor state, for an instruction that the hypervisor does not emulate, the hypervisor should pass control to the operating system at the operating system's Program interrupt vector location, with all registers (SRR0, SRR1, MSR, GPRs, etc.) set as if the instruction had caused a Privileged Instruction type Program interrupt, except with SRR144:45 set to Ob10. (The Hypervisor Emulation Assistance interrupt was added to the architecture in V. 2.05, and the Illegal Instruction type Program interrupt was removed from the architecture in V. 2.06. In V. 2.05 the Hypervisor Emulation Assistance interrupt was optional: implementations that supported it generated it as described in V. 2.06, and never generated an Illegal Instruction type Program interrupt; implementations that did not support it generated an Illegal Instruction type Program interrupt as described above.)


#### Abstract

Programming Note When LPCR $_{\text {EVIRT }}=1$, some of the conditions that cause a Privileged Instruction type Program interrupt when $\mathrm{LPCR}_{\text {EVIRT }}=0$ (attempted execution, in privileged but non-hypervisor state, of a hypervisor privileged instruction or of an mtspr or mfspr instruction specifying an SPR that is hypervisor privileged for the operation) instead cause a Hypervisor Emulation Assistance interrupt. Having these cases cause a Hypervisor Emulation Assistance interrupt permits support of nested hypervisors through virtualization of hypervisor facilities, and simplifies creation of a common kernel for the OS and the hypervisor. Some operating systems may still have code to handle these conditions, at the Program interrupt vector location. For this reason, if a Hypervisor Emulation Assistance interrupt occurs with HSRR1 ${ }_{45}=1$ and the hypervisor is not providing either of these functions, the hypervisor should pass control to the operating system at the operating system's Program interrupt vector location, with all registers (SRR0, SRR1, MSR, GPRs, etc.) set as if the instruction had caused a Privileged Instruction type Program interrupt, including setting $\mathrm{SRR1}_{3} 49$ to $0 \mathrm{b00}$.


### 6.5.10 Floating-Point Unavailable Interrupt

A Floating-Point Unavailable interrupt occurs when no higher priority exception exists, an attempt is made to execute a floating-point instruction (including float-ing-point loads, stores, and moves), and $M S R_{F P}=0$.

The following registers are set:
SRRO Set to the effective address of the instruction that caused the interrupt.
SRR1
33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0800, possibly offset as specified in Figure 65.

### 6.5.11 Decrementer Interrupt

A Decrementer interrupt occurs when no higher priority exception exists, a Decrementer exception exists, and $M S R_{E E}=1$.
The following registers are set:
SRRO Set to the effective address of the instruction that the thread would have attempted
to execute next if no interrupt conditions were present.

## SRR1

33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0900, possibly offset as specified in Figure 65.

### 6.5.12 Hypervisor Decrementer Interrupt

A Hypervisor Decrementer interrupt occurs when no higher priority exception exists, a Hypervisor Decrementer exception exists, and the value of the following expression is 1 .
$\left(\mathrm{MSR}_{\text {EE }}\left|\neg\left(\mathrm{MSR}_{\mathrm{HV}}\right)\right| \mathrm{MSR}_{\text {PR }}\right)$ \& HDICE
The following registers are set:
HSRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present.
HSRR1
33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0980, possibly offset as specified in Figure 65.

## Programming Note

Because the value of MSR ${ }_{\text {EE }}$ is always 1 when the thread is in problem state, the simpler expression

$$
\left(\mathrm{MSR}_{\text {EE }} \mid \neg\left(\mathrm{MSR}_{\mathrm{HV}}\right)\right) \& \text { HDICE }
$$

is equivalent to the expression given above.

### 6.5.13 Directed Privileged Doorbell Interrupt

A Directed Privileged Doorbell interrupt occurs when no higher priority exception exists, a Directed Privileged Doorbell exception is present, and $\mathrm{MSR}_{\mathrm{EE}}=1$. Directed Privileged Doorbell exceptions are generated when Directed Privileged Doorbell messages (see Chapter 10) are received and accepted by the thread.

The following registers are set:
SRRO

Set to the effective address of the instruction that the thread would have attempted
to execute next if no interrupt conditions were present.

## SRR1

33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.

MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0A00, possibly offset as specified in Figure 65.

### 6.5.14 System Call Interrupt

A System Call interrupt occurs when a System Call instruction is executed.

The following registers are set:
SRRO Set to the effective address of the instruction following the System Call instruction.

SRR1
33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0C00, possibly offset as specified in Figure 65.

## Programming Note

An attempt to execute an sce instruction with LEV=1 in problem state should be treated as a programming error.

### 6.5.15 Trace Interrupt

A Trace interrupt occurs when no higher priority exception exists and either a transaction is completed or any instruction except rfid, hrfid, rfscv, or a Power-Saving Mode instruction is successfully completed, provided any of the following is true:

- the instruction is mtmsr[d] and $\mathrm{MSR}_{\mathrm{TE}}=0 \mathrm{~b} 10$ when the instruction was initiated,
- the instruction is not mtmsr[d] and $\mathrm{MSR}_{\mathrm{TE}}=0 \mathrm{~b} 10$,
- the instruction is a Branch instruction and $\mathrm{MSR}_{\text {TE }}=0 \mathrm{b01}$,
- the transaction is completed when $M_{\text {MR }}=0 b 11$, or
- a CIABR match occurs.

Completion of a transaction, for the purposes of Transaction Completion Tracing, means that either the trans-
action has succeeded and execution of the tend. instruction has successfully completed or that failure handling (see Section 5.3.3 of Book II) has completed for any cause of failure other than execution of treclaim.

Successful completion for an instruction means that the instruction caused no other interrupt and, if the thread is in Transactional state, did not cause the transaction to fail in such a way that the instruction did not complete. (See Section 5.3.1 of Book II). Thus a Trace interrupt never occurs for a System Call or System Call Vectored instruction, or for a Trap instruction that traps, or for a dcbf that is executed in Transactional state. The instruction that causes a Trace interrupt is called the "traced instruction".

The following registers are set:
SRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present.

## SRR1

33 Set to 1.
34 Set to 0.
35 Set to 1 if the the Trace interrupt is not the result of a CIABR match and the traced instruction is a Load instruction or is specified to be treated as a Load instruction; otherwise set to 0 .
36 Set to 1 if the the Trace interrupt is not the result of a CIABR match and the traced instruction is a Store instruction or is specified to be treated as a Store instruction; otherwise set to 0 .
43 Set to 1 if the traced instruction is the result of a CIABR match.
44 Set to 1 if the interrupt ocurred as the result of transaction completion.
45:47 Set to 0.
Others Loaded from the MSR.

- Programming Note

Bit 33 is set to 1 for historical reasons.
SIAR For all Trace interrupts other than those caused by a transaction completion or a CIABR match, set to the effective address of the traced instruction; otherwise undefined.

SDAR For all Trace interrupts other than those caused by a transaction completion or a CIABR match, set to the effective address of the storage operand (if any) of the traced instruction; otherwise undefined.

If the state of the Performance Monitor is such that the Performance Monitor may be altering the SIAR and SDAR (i.e., if $M M C R 0_{\text {PMAE }}=1$ ), the contents of the

SIAR and SDAR are undefined for the Trace interrupt and may change even when no Trace interrupt occurs.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_00D0, possibly offset as specified in Figure 65. For a Trace interrupt resulting from execution of an instruction that modifies the value of $\mathrm{MSR}_{\mathrm{IR}}$ or $\mathrm{MSR}_{\mathrm{DR}}$, the Trace interrupt vector location is based on the modified values.

## Programming Note

The following instructions are not traced.

- rfid
- hrfid
- rfscv
- sc, scv, and Trap instructions that trap
- Power-Saving Mode instructions
- other instructions that cause interrupts (other than Trace interrupts)
- the first instructions of any interrupt handler
- instructions that are emulated by software
- instructions, executed in Transactional state, that are disallowed in Transactional state
- instructions, executed in Transactional state, that cause types of accesses that are disallowed in Transactional state
- mtspr, executed in Transactional state, specifying an SPR that is not part of the Transactional Memory checkpointed registers
- tbegin. executed at maximum nesting depth

In general, interrupt handlers can achieve the effect of tracing these instructions.

### 6.5.16 Hypervisor Data Storage Interrupt

A Hypervisor Data Storage interrupt occurs when no higher priority exception exists, the thread is not in hypervisor state, and either
(a) $\mathrm{HR}\left\|\| \mathrm{GR}=0 \mathrm{~b} 00, \mathrm{VPM}_{1}=0, \mathrm{LPCR}_{\mathrm{KBV}}=1\right.$, and a Virtual Storage Page Class Key Protection exception exists or
(b) $\mathrm{HR} \| \mathrm{GR}=0 \mathrm{~b} 00$, the value of the expression
$\left(\neg \mathrm{MSR}_{\mathrm{DR}}\right) \mid\left(\mathrm{VPM}_{1} \& \mathrm{PRTE}_{V} \& \mathrm{MSR}_{\mathrm{DR}}\right)$
is 1 , and a data access cannot be performed for any of the following reasons, or
no higher priority exception exists and either
(c) $\mathrm{HR} \| \mathrm{GR}=0 \mathrm{~b} 00$ and a reference or change bit update cannot be performed as
described below, or
(d) HRIIGR $\neq 0 \mathrm{O} 00$ and partition-scoped translation
prevents an access from being performed for any of the following reasons.
(In the expression for (b) above, "PRTE ${ }_{V}$ " is shorthand indicating that an invalid segment table descriptor did not stop the translation process. Note that an SLB hit may satisfy this condition even when the Process Table Entry is invalid.)
I - HR\|IGR=0b00, data address translation is enabled ( $\mathrm{MSR}_{\mathrm{DR}}=1$ ) and the virtual address of any byte of the storage location specified by a Load, Store, icbi, dcbz, dcbst, or dcbf[ $/]$ instruction cannot be translated to a real address because no valid PTE was found for the VPM translation.

- HRIIGR $\neq 0 b 00$ and the virtual / guest real address of any byte of the storage location specified by a Load, Store, icbi, dcbz, dcbst, or dcbf[I] instruction cannot be translated to a host real address because no valid PTE was found in the parti-tion-scoped page table.
■ The virtual / guest real address of a page directory entry or process table entry could not be translated when HR\|GR $\neq 0 \mathrm{~b} 00$; or the virtual address of a process table entry or segment table entry group could not be translated when $\mathrm{VPM}_{1}=1$ and HR\|GR=0b00.
- An unsupported MMU configuration is found. In addition to an invalid radix tree configuration found in the partition-scoped tables, this type of exception will also be reported outside of hypervisor real mode for translation mode mismatches including $\mathrm{GR}=\mathrm{HR}, \mathrm{UPRT}=0$ when $\mathrm{GR}=1$ or $\mathrm{HR}=1$, LPID=0 if $\mathrm{MSR}_{\mathrm{HV}}=0$ when $\mathrm{GR}=1$ or $\mathrm{HR}=1$, and $\mathrm{HR}=0$ for LPID=0 when HR=1 for another partition ID.
- A reference or change bit update in a parti-tion-scoped PTE cannot be performed (including for the process-scoped PDE or PTE or process table entry for a radix guest or the process table entry or segment table entry group for a paravirtualized HPT guest).


## Programming Note

When reporting failure to set a reference or change bit for a table entry, whether the change bit must be set is inferred from whether the access is reported to be a store. (A load may report store if, when attempting to set the reference bit, the update of the change bit in the partition-scoped PTE mapping the pro-cess-scoped PTE fails.) Behavior is similar for access authority failures.

- HRIIGR=0b00, data address translation is disabled ( $\mathrm{MSR}_{\mathrm{DR}}=0$ ), and the virtual address of any byte of the storage location specified by a Load, Store, icbi, dcbz, dcbst, or dcbf[l] instruction cannot be
translated to a real address by means of the virtual real addressing mechanism.
- The effective address specified by a Iq, stq, Iwat, Idat, Ibarx, Iharx, Iwarx, Idarx, Iqarx, stwat, stdat, stbcx., sthcx., stwcx., stdcx., or stqcx. instruction refers to storage that is Write Through Required or Caching Inhibited; or the effective address specified by a Iwat, Idat, stwat, or stdat instruction refers to storage that is Guarded.
- The effective address specified by a copy or paste[.] instruction refers to storage that is Caching Inhibited.
- The effective address specified by an instruction other than copy or paste[.] refers to CSM.
- An accelerator is specified as the source of a copy instruction or an attempt is made to access an accelerator that is not properly initialized for the software's use.
- A transfer specified to access CSM does not specify (local) main storage as the other end of the transfer.
- The specified CSM address experienced an exception in the outboard translation process.
- The access violates storage protection. In addition to the legacy VPM cases, this includes mismatches in access authority in which the pro-cess-scoped PTE permits the access but the partition-scoped PTE does not. It also includes lack of necessary authority for accesses to pro-cess-scoped tables, for example lack of write authority to set a reference bit in the pro-cess-scoped PTE. (In such a case, the "access" reported as failing would be the access to the pro-cess-scoped table. The HDAR would provide the guest real / (abbreviated) virtual address of the table entry.)
- A Data Address Watchpoint match occurs.

A Hypervisor Data Storage interrupt also occurs when no higher priority exception exists and an attempt is made to execute a Load Atomic or Store Atomic instruction specifying an invalid function code.

If a stbcx., sthcx., stwcx., stdcx., or stqcx. would not perform its store in the absence of a Hypervisor Data Storage interrupt, and either (a) the specified effective address refers to storage that is Write Through Required or Caching Inhibited, or (b) a non-conditional Store to the specified effective address would cause a Hypervisor Data Storage interrupt, it is implementa-tion-dependent whether a Hypervisor Data Storage interrupt occurs.
If the XER specifies a length of zero for an indexed Move Assist instruction, a Hypervisor Data Storage interrupt does not occur.
The following registers are set:
HSRRO Set to the effective address of the instruction that caused the interrupt.
HSRR1

33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.

## MSR See Figure 64.

HDSISR
32
33 Set to 1 if the translation for an attempted access is not found in the Page Table; otherwise set to 0 .
34:35 Set to 0.
36 Set to 1 if the access is not permitted by Figure 43 45, or the privilege, read, or read/write bits in Figure 44 as appropriate; otherwise set to 0 .
37 Set to 1 if the access is due to a Iq, stq, Iwat, Idat, Ibarx, Iharx, Iwarx, Idarx, Iqarx, stwat, stdat, stbcx., sthcx., stwcx., stdcx., or stqcx. instruction that addresses storage that is Write Through Required or Caching Inhibited; or if the access is due to a Iwat, Idat, stwat, or stdat instruction that addresses storage that is Guarded; or if the access is due to a copy or paste[.] instruction that addreses storage that is caching inhibited; or if the access is due to an instruction other than copy or paste[.] addressing CSM; otherwise set to 0 .
38 Set to 1 by an explicit access for a Store or dcbz instruction; set to one when a pro-cess-scoped PTE update fails due to a lack of write authority or the inability to set the change bit in the partition-scoped PTE; otherwise set to 0 .
39:40 Set to 0 .
41 Set to 1 if a Data Address Watchpoint match occurs; otherwise set to 0 .
42 Set to 1 if the access is not permitted by virtual page class key protection; otherwise set to 0 .

43 Set to 0.
44 Set to 1 if an unsupported MMU configuration is found during the translation process. Set to 1 if an attempt to atomically set a reference or change bit fails; otherwise set to 0.

46 Set to 1 if HR\|GR $\neq 0$ b00 and the virtual / guest real address of a page directory entry, page table entry, or process table entry could not be translated; or HR\|GR=0b00, VPM $=1$, and the virtual address of a process table entry or segment table entry group could not be translated; otherwise set to 0 .
47:57 Set to 0.
58 Set to 1 if the specified CSM address experienced an exception in the outboard translation process; otherwise set to 0 .

> Programming Note
> The exception identified by bit 58 may be a problem the hypervisor can fix. Additional information may be retained by the platform for this exception.

59 Set to 1 if a transfer specified to access CSM does not specify (local) main storage as the other end of the transfer; otherwise set to 0 .
60 Set to 1 if an accelerator is specified as the source of a copy instruction or an attempt is made to access an accelerator that is not properly initialized for the software's use; otherwise set to 0 .
61 Set to 1 if an attempt is made to execute a Load Atomic or Store Atomic instruction specifying an invalid function code; otherwise set to 0 .
Set to 0 .
HDAR Set to the effective address or portion of the VPN of a storage element, or undefined, as described in the following list. The list should be read from the top down; the HDAR is set as described by the first item that corresponds to an exception that is reported in the HDSISR. For example, if a Load Word instruction causes a storage protection violation and a Data Address Watchpoint match (and both are reported in the HDSISR), the HDAR is set to the effective address of a byte in the first aligned doubleword for which access was attempted in the page that caused the exception.

- least significant 64 bits of the VA of the table entry or group when a process table entry or segment table entry group virtual address cannot be translated in Paravirtualized HPT mode with $V_{P M}=1$.
- EA, when a Hypervisor Data Storage exception occurs for reasons other than a Data Address Watchpoint match
- a byte in the block that caused the exception, for a Cache Management instruction
- a byte in the first aligned quadword for which access was attempted in the page that caused the exception, for a quadword Load or Store instruction (i.e., a Load or Store instruction for which the storage operand is a quadword; "first" refers to address order: see Section 6.7)
- a byte in the first aligned doubleword for which access was attempted in the page that caused the exception, for a non-quadword Load or Store instruction
- undefined, for a Data Address Watchpoint match

For the cases in which the HDAR is specified above to be set to an effective address, if the interrupt occurs in 32-bit mode the high-order 32 bits of the HDAR are set to 0 .

## Programming Note

Note that for HPT translation, the full EA is a superset of the bits required to construct the full VA, when also provided with the VSID in the ASDR.

ASDR When HR\|GR=0b00, loaded with VSID, B, Ks, Kp, N, C, L, and LP values from the segment descriptor that translated the access or indicated the base of the table, or undefined, as described in the following list. For a large segment the values of the bits below the VSID are undefined. When HR\|GR $=0$ b00 (nested translaiton is taking place), loaded with the guest real address down to bit 51 of a storage element or table entry, or undefined, as described in the following list. (Note that the size of the GRA may differ depending on whether the host uses HPT or radix tree translation. The GRA is effectively an abbreviated VSID for an HPT host, while its size is determined by the maximum size of a radix tree for a guest that uses radix tree translation.) The list should be read from the top down; the ASDR is set as described by the first item that corresponds to an exception that is reported in the HDSISR.

- the guest real page address of the table entry when a process table or process-scoped page directory or page table entry guest real address cannot be translated or the VSID of the table entry when a process or segment table entry virtual address cannot be translated (the rest of the segment descriptor is implied).
■ the guest real address of the pro-cess-scoped PDE or PTE or process table entry when a reference or change bit in the partition-scoped PTE mapping the process-scoped PDE or PTE or process table entry cannot be set atomically
- the guest real address of the storage element when a reference or change bit in the partition-scoped PTE cannot be set atomically
- the guest real address of the storage element, process table entry, page directory entry, or page table entry (depending on which partition-scoped table has the flaw) for an unsupported radix tree configuration in the parti-tion-scoped table (the effective address for other cases of the invalid MMU configuration exception is found in the HDAR)
- the guest real address of the pro-cess-scoped PTE when an attempt is made to set a reference or change bit without write authority in the parti-tion-scoped PTE that maps it
- the guest real address or segment descriptor associated with the specified storage element when a Hypervisor Data Storage exception occurs for reasons other than a Data Address Watchpoint match
- undefined, for a Data Address Watchpoint match, unsupported MMU configuration, or accesses to storage that is Caching Inhibited or Write Through Required by the instructions that are prohibited from making such accesses.

If multiple Hypervisor Data Storage exceptions occur for a given effective address, any one or more of the bits corresponding to these exceptions may be set to 1 in the HDSISR. If the HDSISR reports other exceptions together with a Virtualized Page Class Key Storage Protection exception that occurs when LPCR KBV $=1$ and Virtualized Partition Memory is disabled by VPM $_{1}=0$, the other exceptions are actually DSIs.

## Programming Note

A Virtual Page Class Key Storage Protection exception that occurs with LPCR ${ }_{\text {KBV }}=1$ and Virtualized Partition Memory disabled by $\mathrm{VPM}_{1}=0$ identifies an access that must be emulated by the hypervisor. When it is reported together with other exceptions in the HDSISR, the hypervisor should service the Virtual Page Class Key Storage Protection exception first. This is in part because the operating system may be using some PTE fields for non-architected purposes, which could in turn cause spurious exceptions to be reported.

Execution resumes at effective address 0x0000_0000_0000_0E00, possibly offset as specified in Figure 65.

### 6.5.17 Hypervisor Instruction Storage Interrupt

A Hypervisor Instruction Storage interrupt occurs when the thread is not in hypervisor state, no higher priority exception exists, and either
(a) $\mathrm{HR} \| G R=0 b 00$, the value of the expression

$$
\left.\left(\neg \mathrm{MSR}_{\mathrm{IR}}\right) \text { I }\left(\mathrm{VPM}_{1} \& \mathrm{PRTE}_{\mathrm{V}} \& \mathrm{MSR}_{\mathrm{IR}}\right)\right)
$$

is 1 , and the next instruction to be executed cannot be fetched for any of the following reasons, or
(b) HRIIGR $\neq 0 \mathrm{bOO}$ and partition-scoped translation prevents the next instruction to be executed from being fetched for any of the following reasons.
(In the expression for (a) above, "PRTE $\mathrm{V}_{\mathrm{V}}$ " is shorthand indicating that an invalid segment table descriptor did not stop the translation process. Note that an SLB hit may satisfy this condition even when the Process Table Entry is invalid.)

A Hypervisor Instruction Storage interrupt also occurs when no higher priority exception exists, HR\|GR=0b00, and a reference or change bit update cannot be performed as

## described below.

■ Instruction address translation is enabled ( $\mathrm{MSR}_{\mathrm{IR}^{2}}=1$ ) and the virtual address cannot be translated to a real address because no valid PTE was found for the VPM translation.

- HRIIGR $\neq 0$ b00 and the guest real address of the instruction cannot be translated to a host real address because no valid PTE was found in the partition-scoped page table.
- The virtual / guest real address of a page directory entry or process table entry could not be translated when HRIIGR $\neq 0 \mathrm{~b} 00$; or the virtual address of a process table entry or segment table entry group could not be translated when $\mathrm{VPM}_{1}=1$ and HR\|GR=0b00.
- An unsupported MMU configuration is found. In addition to an invalid radix tree configuration found in the partition-scoped tables, this type of exception will also be reported outside of hypervisor real mode for translation mode mismatches including $G R \neq H R, U P R T=0$ when $G R=1$ or $H R=1, L P I D=0$ if $\mathrm{MSR}_{\mathrm{HV}}=0$ when $\mathrm{GR}=1$ or $\mathrm{HR}=1$, and $\mathrm{HR}=0$ for LPID=0 when HR=1 for another partition ID.
- A reference or change bit update in a parti-tion-scoped PTE cannot be performed (including for the process-scoped PDE or PTE or process table entry for a radix guest or the process table entry or segment table entry group for a paravirtualized HPT guest).
- HRIIGR=0b00, instruction address translation is disabled $\left(\mathrm{MSR}_{\mathrm{IR}}=0\right)$, and the virtual address can-
not be translated to a real address by means of the virtual real addressing mechanism.
■ The fetch violates storage protection. In addition to the legacy VPM cases, this includes mismatches in access authority in which the pro-cess-scoped PTE permits the access but the partition-scoped PTE does not. It also includes lack of necessary authority for accesses to pro-cess-scoped tables, for example lack of write authority to set a reference bit in the pro-cess-scoped PTE. (In such a case, the "access" reported as failing would be the access to the pro-cess-scoped table. The HDAR would provide the guest real / (abbreviated) virtual address of the table entry.)

The following registers are set:
HSRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present (if the interrupt occurs on attempting to fetch a branch target, HSRRO is set to the branch target address).

## HSRR1

I

34 Set to 0.
35 Set to 1 if the access is to No-execute (as indicated by the N bit in the segment table entry or the N bit in the HPT PTE or the Execute and Privilege bits in the EAA field of the Radix PTE and IAMR key 0 ) or Guarded storage; otherwise set to 0 .
36 Set to 1 if the access is not permitted by Figure 43 45, or the privilege, read, or write bits in Figure 44 as appropriate; otherwise set to 0 .
42 Set to 1 if the access is not permitted by virtual page class key protection; otherwise set to 0 .
43 Set to 0.
44 Set to 1 if an unsupported MMU configuration is found during the translation process.
45 Set to 1 if an attempt to atomically set a reference or change bit fails; otherwise set to 0.

46 Set to 1 if HR\|GR $\neq 0 \mathrm{bOO}$ and the virtual / guest real address of a page directory entry, page table entry, or process table entry could not be translated; or HR\|GR=0b00, VPM $=1$, and the virtual address of a process table entry or segment table entry group could not be translated; otherwise set to 0 .
47 Set to 1 if the operation that caused the exception was attempting to update storage; otherwise set to 0 . This bit may be set as a modifier to bit 45 to indicate that a
change bit must be set. It may also be set as a modifier to bits 36 and 42, to indicate that write authority was required to complete the operation.
Others Loaded from the MSR.
HDAR Set to the least significant 64 bits of the VA of a table entry or group when HR\|GR=0b00 and a process table entry or segment table entry group virtual address cannot be translated and $\mathrm{VPM}_{1}=1$. May be set spuriously in other cases.
ASDR When HR\|GR=0b00, loaded with VSID, B, $\mathrm{Ks}, \mathrm{Kp}, \mathrm{N}, \mathrm{C}, \mathrm{L}$, and LP values from the segment descriptor that translated the access or indicated the base of the table, or undefined, as described in the following list. For a large segment the values of the bits below the VSID are undefined. When HR\|GR $\neq 0 b 00$ (nested translaiton is taking place), set to the guest real address down to bit 51 of the instruction or table entry, or undefined, as described in the following list.

- the guest real address of the table entry when a process table or pro-cess-scoped page directory or page table entry guest real address cannot be translated or the VSID of the table entry when a process or segment table entry virtual address cannot be translated (the rest of the segment desrcriptor is implied).
- the guest real address of the pro-cess-scoped PDE or PTE or process table entry when a reference or change bit in the partition-scoped PTE mapping the process-scoped PDE or PTE or process table entry cannot be set atomically
- the guest real address of the instruction when a reference or change bit in the partition-scoped PTE cannot be set atomically
■ the guest real address of the instruction, process table entry, page directory entry, or page table entry (depending on which partition-scoped table has the flaw) for an unsupported radix tree configuration in the parti-tion-scoped table (the effective address for other cases of the invalid MMU configuration exception will be found in HSRRO)
■ the guest real address of the pro-cess-scoped PTE when an attempt is made to set a reference bit without write authority in the partition-scoped PTE that maps it
- the guest real address or segment descriptor associated with the instruction that the thread would have attempted to execute next if no interrupt conditions were present (parti-tion-scoped page fault or protection exception)
- undefined for unsupported MMU configuration
MSR See Figure 64.
If multiple Hypervisor Instruction Storage exceptions occur due to attempting to fetch a single instruction, any one or more of the bits corresponding to these exceptions may be set to 1 in HSRR1.

Execution resumes at effective address 0x0000_0000_0000_0E10, possibly offset as specified in Figure 65.

### 6.5.18 Hypervisor Emulation Assistance Interrupt

A Hypervisor Emulation Assistance interrupt is generated when execution is attempted of an illegal instruction, or of a reserved instruction or an instruction that is not provided by the implementation. It is also generated under the following conditions.
■ an attempt to execute a hypervisor privileged instruction when $\mathrm{MSR}_{\text {HVIIPR }}=0 \mathrm{~b} 00$ and $\mathrm{LPCR}_{\mathrm{E}}$ VIRT $=1$

- an attempt to execute an mtspr or mfspr instruction that specifies an SPR that is hypervisor privileged for the operation when $\mathrm{MSR}_{\mathrm{HVIIPR}}=0 \mathrm{bOO}$ and LPCR ${ }_{\text {EVIRT }}=1$
- an mtspr or mfspr instruction is executed when $M S R_{P R}=1$ if the instruction specifies an SPR with $\mathrm{spr}_{0}=0$ that is not provided by the implementation
- an mtspr or mfspr instruction is executed when $M_{R R}=0$ if the instruction specifies SPR 0
- an mfspr instruction is executed when $M S R_{P R}=0$ if the instruction specifies SPR 4, 5, or 6
■ an mtspr or mfspr instruction is executed when $M_{S R}=0$ and $L_{P C R}^{E V I R T}=1$ if the instruction specifies an SPR other than those listed above that is not provided by the implementation

A Hypervisor Emulation Assistance interrupt may be generated when execution is attempted of an instruction that is in invalid form or that is treated as if the instruction form were invalid.

The following registers are set:
HSRRO Set to the effective address of the instruction that caused the interrupt.

## HSRR1

33:36
Set to 0 .
Set to 0 .
Set to 1 for an attempt, when $\mathrm{MSR}_{\text {HV PR }}=$ ObOO and LPCR ${ }_{\text {EVIRT }}=1$, to execute a
hypervisor privileged instruction or an $\boldsymbol{m t s p r}$ or mfspr instruction that specifies an SPR that is hypervisor privileged for the operation; otherwise set to 0 .
46:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
HEIR Set to a copy of the instruction that caused the interrupt
Execution resumes at effective address 0x0000_0000_0000_0E40, possibly offset as specified in Figure 65.

## Programming Note

If a Hypervisor Emulation Assistance interrupt occurs with HSRR1 ${ }_{45}=0$ when the thread is not in hypervisor state, for an instruction that the hypervisor does not emulate, the hypervisor should pass control to the operating system as if the instruction had caused an "Illegal Instruction type Program interrupt", as described in a Programming Note near the end of Section 6.5.9, "Program Interrupt" on page 1071.

Similarly, if a Hypervisor Emulation Assistance interrupt occurs with HSRR1 ${ }_{45}=1$ when the thread is in privileged non-hypervisor state, for an instruction that the hypervisor does not virtualize, the hypervisor should pass control to the operating system as if the instruction had caused a Privileged Instruction type Program interrupt, as described in another Programming Note near the end of Section 6.5.9, "Program Interrupt" on page 1071.

## Programming Note

In versions of the architecture that precede V 3.0, the policy implemented directly in hardware was that an attempt when $M_{S R}=0$ to execute an $\boldsymbol{m t s p r}$ or mfspr instruction specifying an SPR that was not implemented (with the exception of SPR 0 for $\boldsymbol{m t s p r}$ and SPRs $0,4,5$, and 6 for $\boldsymbol{m f s p r}$ ) would be treated as a noop. These former noop cases now cause a Hypervisor Emulation Assistance interrupt when $\mathrm{LPCR}_{\text {EVIRT }}=1$ to enable future functions to be emulated on older implementations. If there is no future function emulation to be performed, hypervisor software must choose a policy from the following.
■ treat the instruction as an error

- implement the legacy noop behavior directly (emulate the old behavior)
■ give control to the operating system


### 6.5.19 Hypervisor Maintenance Interrupt

A Hypervisor Maintenance interrupt occurs when no higher priority exception exists, a Hypervisor Maintenance exception exists (a bit in the HMER is set to one), the exception is enabled in the HMEER, and the value of the following expression is 1 .
$\left(M S R_{E E}\left|\neg\left(\mathrm{MSR}_{H V}\right)\right| \mathrm{MSR}_{P R}\right)$
The following registers are set:
HSRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present.

HSRR1
33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
HMER See Section 6.2.9 on page 1049.

The exception bits in the HMER are sticky; that is, once set to 1 they remain set to 1 until they are set to 0 by an mthmer instruction.

Execution resumes at effective address 0x0000_0000_0000_0E60.

## Programming Note

Because the value of MSR ${ }_{E E}$ is always 1 when the thread is in problem state, the simpler expression

$$
\left(\mathrm{MSR}_{\mathrm{EE}} \mid \neg\left(\mathrm{MSR}_{\mathrm{HV}}\right)\right)
$$

is equivalent to the expression given above.

## Programming Note

If an implementation uses the HMER to record that a readable resource, such as the Time Base, has been corrupted, then, because the HMI is disabled in the hypervisor state, it is necessary for the hypervisor to check HMER after reading that resource to be sure an error has not occurred.

### 6.5.20 Directed Hypervisor Doorbell Interrupt

A Directed Hypervisor Doorbell interrupt occurs when no higher priority exception exists, a Directed Hypervisor Doorbell exception is present, and the value of the following expression is 1 .
$\left(\mathrm{MSR}_{\mathrm{EE}}\left|\neg\left(\mathrm{MSR}_{\mathrm{HV}}\right)\right| \mathrm{MSR}_{\mathrm{PR}}\right)$

Directed Hypervisor Doorbell exceptions are generated when Directed Hypervisor Doorbell messages (see Chapter 10) are received and accepted by the thread.
The following registers are set:
HSRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present.

## HSRR1

33:36 Set to 0 .
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0E80, possibly offset as specified in Figure 65.

## Programming Note

Because the value of MSR $_{\text {EE }}$ is always 1 when the thread is in problem state, the simpler expression

$$
\left(\mathrm{MSR}_{\mathrm{EE}} \mid \neg\left(\mathrm{MSR}_{\mathrm{HV}}\right)\right)
$$

is equivalent to the expression given above.

### 6.5.21 Hypervisor Virtualization Interrupt

A Hypervisor Virtualization interrupt occurs when no higher priority exception exists, a Hypervisor Virtualization exception exists, and the value of the following equation is1.
$\left(\mathrm{MSR}_{\text {EE }}\left|\neg\left(\mathrm{MSR}_{\text {HV }}\right)\right| \mathrm{MSR}_{\text {PR }}\right) \&$ HVICE
The occurrence of the interrupt does not cause the exception to cease to exist.
HSRRO Set to the effective address of the instruction that the thread would have attempted to execute next if no interrupt conditions were present.

## HSRR1

33:36 Set to 0 .
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0EA0, possibly offset as specified in Figure 65.

### 6.5.22 Performance Monitor Interrupt

A Performance Monitor interrupt occurs when no higher priority exception exists, a Performance Monitor exception exists, event-based branches are disabled $\left(M M C R 0_{E B E}=0\right)$, and $M S R_{E E}=1$, and either $\mathrm{HFSCR}_{\mathrm{PM}}=1$ or the thread is in hypervisor state.

If multiple Performance Monitor exceptions occur before the first causes a Performance Monitor interrupt, the interrupt reflects the most recent Performance Monitor exception and the preceding Performance Monitor exceptions are lost.

The following registers are set:
SRRO Set to the effective address of the instruction that would have been attempted to be execute next if no interrupt conditions were present.

## SRR1

33:36 and 42:47
Reserved.
Others Loaded from the MSR.

MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0F00, possibly offset as specified in Figure 65.

### 6.5.23 Vector Unavailable Interrupt

A Vector Unavailable interrupt occurs when no higher priority exception exists, an attempt is made to execute a Vector instruction (including Vector loads, stores, and moves), and MSR ${ }_{\text {VEC }}=0$.
The following registers are set:
SRRO Set to the effective address of the instruction that caused the interrupt.
SRR1
33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0F20, possibly offset as specified in Figure 65.

### 6.5.24 VSX Unavailable Interrupt

A VSX Unavailable interrupt occurs when no higher priority exception exists, an attempt is made to execute a VSX instruction (including VSX loads, stores, and moves), and $\mathrm{MSR}_{\mathrm{VSX}}=0$.

The following registers are set:
SRRO Set to the effective address of the instruction that caused the interrupt.

SRR1
33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at effective address 0x0000_0000_0000_0F40, possibly offset as specified in Figure 65.

### 6.5.25 Facility Unavailable Interrupt

A Facility Unavailable interrupt occurs when no higher priority exception exists, and one of the following occurs.

- a facility is accessed in problem state when it has been made unavailable by the FSCR
- a Performance Monitor register is accessed or a clrbhrb or mfbhrbe instruction is executed in problem state when it has been made unavailable by MMCRO.
- the Transactional Memory Facility is accessed in any privilege state when it has been made unavailable by $\mathrm{MSR}_{\text {TM }}$.

The following registers are set:
SRRO Set to the effective address of the instruction that caused the interrupt.

SRR1
33:36 Set to 0.
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
FSCR
0:7 See Section 6.2.11 on page 1049.
Others Not changed.
Execution resumes at effective address 0x0000_0000_0000_0F60, possibly offset as specified in Figure 65.

## Programming Note

For the case of an outer tbegin., the interrupt handler should either return to the tbegin. with $\mathrm{MSR}_{\mathrm{TM}}$ $=1$ (allowing the program to use transactions), or treat the attempt to initiate an outer transaction as a program error.

### 6.5.26 Hypervisor Facility Unavailable Interrupt

A Hypervisor Facility Unavailable interrupt occurs when no higher priority exception exists, and one of the following occurs.

- a facility is accessed in problem or privileged non-hypervisor states when it has been made unavailable by the HFSCR.
- The stop instruction is executed in privileged non hypervisor state when any of the following conditions exist.
PSSCR $_{\text {EC }}=1$
PSSCR $_{\text {ESL }}=1$
PSSCR $_{\text {MTL }}>$ PSSCR $_{\text {PSLL }}$
PSSCR $_{\text {RL }}>$ PSSCR $_{\text {PSLL }}$
The following registers are set:
HSRRO Set to the effective address of the instruction that caused the interrupt.
HSRR1
33:36 Set to 0
42:47 Set to 0 .
Others Loaded from the MSR.
MSR See Figure 64 on page 1061.
HFSCR
0:7 See Section 6.2.12 on page 1051.
Others Not changed.
Execution resumes at effective address 0x0000_0000_0000_0F80, possibly offset as specified in Figure 65.


### 6.5.27 System Call Vectored Interrupt

A System Call Vectored interrupt occurs when a SysI tem Call Vectored instruction is executed.

The following registers are set:
LR Set to the effective address of the instruction following the System Call Vectored instruction.

## CTR

33:36 undefined
42:47 undefined
Others Loaded from corresponding bits of the MSR.
MSR See Figure 64 on page 1061.
Execution resumes at the effective address specified in Figure 65

## Programming Note

When the System Call Vectored interrupt sets $\mathrm{MSR}_{\mathrm{IR}}$ to 1, the effective address described above is translated to a real address before being used to access storage. If the effective address cannot be translated, or if instructions cannot be fetched from the addressed storage location (e.g., the access would violate storage protection, or would be to No-execute storage), an Instruction Storage interrupt occurs before the first instruction at the effective address is executed.

Because the System Call Vectored interrupt uses save/restore registers that differ from those used by other interrupts, the System Call Vectored interrupt handler can run with address translation enabled and External interrupts enabled. Similarly, the Programming Note about managing $\mathrm{MSR}_{\text {RI }}$ in s does not apply to the System Call Vectored interrupt handler (the System Call Vectored interrupt does not alter $\mathrm{MSR}_{\mathrm{RI}}$ ).

### 6.6 Partially Executed Instructions

If a Data Storage, Data Segment, Alignment, sys-tem-caused, or imprecise exception occurs while a Load or Store instruction is executing, the instruction may be aborted. In such cases the instruction is not completed, but may have been partially executed in the following respects.

- Some of the bytes of the storage operand may have been accessed, except that if access to a given byte of the storage operand would violate storage protection, that byte is neither copied to a register by a Load instruction nor modified by a Store instruction. Also, the rules for storage accesses given in Section 5.8.1, "Guarded Storage" and in Section 2.2 of Book II are obeyed.
- Some registers may have been altered as described in the Book II section cited above.
- Reference and Change bits may have been updated as described in Section 5.7.13.
■ For a stbcx., sthcx., stwcx., stdcx., or stqcx. instruction that is executed in-order, CRO may have been set to an undefined value and the reservation may have been cleared.

The architecture does not support continuation of an aborted instruction but intends that the aborted instruction be re-executed if appropriate.

## Programming Note

An exception may result in the partial execution of a Load or Store instruction. For example, if the Page Table Entry that translates the address of the storage operand is altered, by a program running on another thread, such that the new contents of the Page Table Entry preclude performing the access, the alteration could cause the Load or Store instruction to be aborted after having been partially executed.

As stated in the Book II section cited above, if an instruction is partially executed the contents of registers are preserved to the extent that the instruction can be re-executed correctly. The consequent preservation is described in the following list. For any given instruction, zero, one, or two items in the list apply.

- For a fixed-point Load instruction that is not a multiple or string form, if $R T=R A$ or $R T=R B$ then the contents of register RT are not altered.
- For an Iq instruction, if $R T+1=R A$ then the contents of register RT+1 are not altered.
- For an update form Load or Store instruction, the contents of register RA are not altered.


### 6.7 Exception Ordering

Since multiple exceptions can exist at the same time and the architecture does not provide for reporting more than one interrupt at a time, the generation of more than one interrupt is prohibited. Some exceptions, such as the Mediated External exception, persist and can be deferred. However, other exceptions would be lost if they were not recognized and handled when they occur. For example, if an External interrupt was generated when a Data Storage exception existed, the Data Storage exception would be lost. If the Data Storage exception was caused by a Store Multiple instruction for which the storage operand crosses a virtual page boundary and the exception was a result of attempting to access the second virtual page, the store could have modified locations in the first virtual page even though it appeared that the Store Multiple instruction was never executed.

For the above reasons, all exceptions are prioritized with respect to other exceptions that may exist at the same instant to prevent the loss of any exception that is not persistent. Some exceptions cannot exist at the same instant as some others.

Data Storage, Hypervisor Data Storage, Data Segment, and Alignment exceptions and transaction failure due to attempted access of a disallowed type while in Transactional state occur as if the storage operand were accessed one byte at a time in order of increasing effective address (with the obvious caveat if the operand includes both the maximum effective address and effective address 0 ). (The required ordering of exceptions on components of non-atomic accesses does not extend to the performing of the component accesses in the event of an exception. For example, if byte $n$ causes a data storage exception, it is not necessarily true that the access to byte $\mathrm{n}-1$ has been performed.)

### 6.7.1 Unordered Exceptions

The exceptions listed here are unordered, meaning that they may occur at any time regardless of the state of the interrupt processing mechanism. These exceptions are recognized and processed when presented.

1. System Reset
2. Machine Check

### 6.7.2 Ordered Exceptions

The exceptions listed here are ordered with respect to the state of the interrupt processing mechanism. With one exception, in the following list the hypervisor forms of the Data Storage and Instruction Storage exceptions can be substituted for the non-hypervisor forms since the hypervisor forms cannot be caused by the same instruction and have the same ordering. The exception
is that Virtual Page Class Key Storage Protection exceptions that occur when $\mathrm{LPCR}_{\mathrm{KBV}}=1$ and Virtualized Partition Memory is disabled by VPM $_{1}=0$ cause only a Hypervisor Data Storage exception (and never a Data Storage exception).

## System-Caused or Imprecise

1. Program

- Imprecise Mode Floating-Point Enabled Exception

2. Hypervisor Maintenance

I 3. Hypervisor Virtualization, External, [Hypervisor] Decrementer, Performance Monitor, Directed Privileged Doorbell, Directed Hypervisor Doorbell

## Instruction-Caused and Precise

1. Instruction Segment
2. [Hypervisor] Instruction Storage
3.a Hypervisor Emulation Assistance
3.b Program

- Privileged Instruction

4. Function-Dependent
4.a Fixed-Point and Branch

1 Hypervisor Facility Unavailable
2 Facility Unavailable
3a Program

- Trap
- TM Bad Thing

3b System Call or System Call Vectored
3c. 1 Data Storage for the case of Fixed-Point Load or Store Caching Inhibited instructions with $M_{S R}=1$ or the case of an invalid function code for an Atomic Memory Operation
3c. 2 all other Data Storage, Hypervisor Data Storage, [Hypervisor] Data Segment, or Alignment
4 Trace
4.b Floating-Point

1 Hypervisor Facility Unavailable
2 FP Unavailable
3a Program - Precise Mode Floating-Pt Enabled Excep'n

3b [Hypervisor] Data Storage, [Hypervisor] Data Segment, or Alignment
4 Trace
4.c Vector

1 Hypervisor Facility Unavailable
2 Vector Unavailable
3a [Hypervisor] Data Storage, [Hypervisor] Data Segment, or Alignment
4 Trace
4.d VSX

1 Hypervisor Facility Unavailable
2 VSX Unavailable
3a Program

- Precise Mode Floating-Pt Enabled Excep'n

3b [Hypervisor] Data Storage, [Hypervisor] Data Segment, or Alignment

4 Trace
4.e Other Instructions

1 Hypervisor Facility Unavailable
2 Facility Unavailable
3a [Hypervisor] Data Storage, [Hypervisor] Data Segment, or Alignment
4 Trace
For implementations that execute multiple instructions in parallel using pipeline or superscalar techniques, or combinations of these, it can be difficult to understand the ordering of exceptions. To understand this ordering it is useful to consider a model in which each instruction is fetched, then decoded, then executed, all before the next instruction is fetched. In this model, the exceptions
a single instruction would generate are in the order shown in the list of instruction-caused exceptions. Exceptions with different numbers have different ordering. Exceptions with the same numbering but different lettering are mutually exclusive and cannot be caused
I by the same instruction. The Hypervisor Virtualization, External, [Hypervisor] Decrementer, Performance Monitor, Directed Privileged Doorbell, and Directed Hypervisor Doorbell interrupts have equal ordering. Similarly, where Data Storage, Data Segment, and Alignment exceptions are listed in the same item they have equal ordering.

Even on threads that are capable of executing several instructions simultaneously, or out of order, instruc-tion-caused interrupts (precise and imprecise) occur in program order.

### 6.8 Event-Based Branch Exception Ordering

Event-based exceptions are not ordered because they can occur simultaneously. Whenever an event-based exception occurs and the exception is enabled, the corresponding "exception occurred" bit in the BESCR is set to 1. See Section 7.2.1 of Book II.

### 6.9 Interrupt Priorities

This section describes the relationship of nonmaskable, maskable, precise, and imprecise interrupts. In the following descriptions, the interrupt mechanism waiting for all possible exceptions to be reported includes only exceptions caused by previously initiated instructions (e.g., it does not include waiting for the Decrementer to step through zero). The exceptions are listed in order of highest to lowest priority. The phrase "corresponding interrupt" means the interrupt having the same name as the exception unless the thread is in power-saving mode, in which case the phrase means the System Reset interrupt.
Unless otherwise stated or obvious from context, it is assumed below that one of the following conditions is satisfied.

■ The thread is not in power-saving mode and the interrupt, unless it is the Machine Check interrupt, is not disabled. (For the Machine Check interrupt no assumption is made regarding enablement.)

- The thread is in power-saving mode and the exception is enabled to cause exit from the mode.
With one exception, in the following list the hypervisor forms of the Data Storage and Instruction Storage exceptions can be substituted for the non-hypervisor forms since the hypervisor forms cannot be caused by
the same instruction and have the same priority. The exception is that exceptions caused by Virtual Page Class Key Storage Protection exceptions that occur when LPCR $_{\text {KBV }}=1$ and Virtualized Partition Memory is disabled by VPM $_{1}=0$ cause only a Hypervisor Data Storage exception (and never a Data Storage exception).

1. System Reset

System Reset exception has the highest priority of all exceptions. If this exception exists, the interrupt mechanism ignores all other exceptions and generates a System Reset interrupt.
Once the System Reset interrupt is generated, no nonmaskable interrupts are generated due to exceptions caused by instructions issued prior to the generation of this interrupt.
2. Machine Check

Machine Check exception is the second highest priority exception. If this exception exists and a System Reset exception does not exist, the interrupt mechanism ignores all other exceptions and generates a Machine Check interrupt.
Once the Machine Check interrupt is generated, no nonmaskable interrupts are generated due to exceptions caused by instructions issued prior to the generation of this interrupt.
3. Instruction-Caused and Precise

This exception is the third highest priority exception. When this exception is created, the interrupt mechanism waits for all possible Imprecise exceptions to be reported. It then generates the appropriate ordered interrupt if no higher priority exception exists when the interrupt is to be generated. Within this category a particular instruction may present more than a single exception. When this occurs, those exceptions are ordered in priority as indicated in the following lists. Where [Hypervisor] Data Storage, Data Segment, and Alignment exceptions are listed in the same item they have equal priority (i.e., the hardware may generate any one of the three interrupts for which an exception exists). For instructions that are forbidden in Transactional state, transaction failure takes priority over all interrupts except Privileged Instruction type Program Interrupts. For data accesses that are forbidden in Transactional state, transaction failure has the same priority as the group of "other" [Hypervisor] Data Storage, Data Segment, and Alignment exceptions. (See Section 5.3.1 of Book II ).
A. Fixed-Point Loads and Stores
a. These exceptions are mutually exclusive and have the same priority:

- Hypervisor Emulation Assistance
- Program - Privileged Instruction
b. Hypervisor Facility Unavailable
c. Facility Unavailable
d. Data Storage for the case of Fixed-Point Load or Store Caching Inhibited instructions with $\mathrm{MSR}_{\mathrm{DR}}=1$ or the case of an invalid function code for an Atomic Memory Operation
e. all other Data Storage, Hypervisor Data Storage, [Hypervisor] Data Segment, or Alignment
f. Trace
B. Floating-Point Loads and Stores
a. Hypervisor Emulation Assistance
b. Hypervisor Facility Unavailable
c. Floating-Point Unavailable
d. [Hypervisor] Data Storage, [Hypervisor] Data Segment, or Alignment
e Trace
C. Vector Loads and Stores
a. Hypervisor Emulation Assistance
b. Hypervisor Facility Unavailable
c. Vector Unavailable
d. [Hypervisor] Data Storage, [Hypervisor] Data Segment, or Alignment
e. Trace
D. VSX Loads and Stores
a. Hypervisor Emulation Assistance
b. Hypervisor Facility Unavailable
c. VSX Unavailable
d. [Hypervisor] Data Storage, [Hypervisor] Data Segment, or Alignment
e. Trace
E. Other Floating-Point Instructions
a. Hypervisor Emulation Assistance
b. Hypervisor Facility Unavailable
c. Floating-Point Unavailable
d. Program - Precise Mode Floating-Point Enabled Exception
e. Trace
F. Other Vector Instructions
a. Hypervisor Emulation Assistance
b. Hypervisor Facility Unavailable
c. Vector Unavailable
d. Trace
G. Other VSX Instructions
a. Hypervisor Emulation Assistance
b. Hypervisor Facility Unavailable
c. VSX Unavailable
d. Program - Precise Mode Floating-Point Enabled Exception
e. Trace
H. TM instruction, mt/fspr specifying TM SPR
a. Program - Privileged Instruction (only for
treclaim., trechkpt.., and mtspr)
b Hypervisor Facility Unavailable
c Facility Unavailable
d Program - TM Bad Thing (only for treclaim.,
trechkpt., and mtspr)
e Trace only mtmsr.
J. Other Instructions
a. Hypervisor Emulation Assistance, for rfscv
b. Program - Privileged Instruction for all except rfebb
c. Hypervisor Facility Unavailable (rfebb only)
d. Facility Unavailable (rfebb only)
e Program - TM Bad Thing for all except
f. Program - Floating-Point Enabled Exception or all except rfebb
g. Trace, for mtmsr[d] and rfebb only
a.These exceptions are mutually exclusive and have the same priority:
■ Program - Trap
- System Call
- System Call Vectored
- Program - Privileged Instruction
- Hypervisor Emulation Assistance
b. Hypervisor Facility Unavailable
c. Facility Unavailable
d. Trace
K. [Hypervisor] Instruction Storage and Instruction Segment
These exceptions have the lowest priority in this category. They are recognized only when all instructions prior to the instruction causing one of these exceptions appear to have completed and that instruction is the next instruction to be executed. The two exceptions are mutually exclusive.
The priority of these exceptions is specified for completeness and to ensure that they are not given more favorable treatment. It is acceptable for an implementation to treat these exceptions as though they had a lower priority.

4. Program - Imprecise Mode Floating-Point Enabled Exception
This exception is the fourth highest priority exception. When this exception is created, the interrupt mechanism waits for all other possible exceptions to be reported. It then generates this interrupt if no higher priority exception exists when the interrupt is to be generated.
5. Hypervisor Maintenance

This exception is the fifth highest priority exception. When this exception is created, the interrupt mechanism waits for all other possible exceptions to be reported. It then generates this interrupt if no higher priority exception exists when the interrupt is to be generated.
If a Hypervisor Maintenance exception exists and each attempt to execute an instruction when the

Hypervisor Maintenance interrupt is enabled causes an exception (see the Programming Note below), the Hypervisor Maintenance interrupt is not delayed indefinitely.
6. Hypervisor Virtualization, Direct External, Mediated External, and [Hypervisor] Decrementer, Performance Monitor, Directed Privileged Doorbell, Directed Hypervisor Doorbell

These exceptions are the lowest priority exceptions. All have equal priority (i.e., the hardware may generate any one of the corresponding interrupts for which an exception exists). When one of these exceptions is created, the interrupt processing mechanism waits for all other possible exceptions to be reported. It then generates the corresponding interrupt if no higher priority exception exists when the interrupt is to be generated.
If a Hypervisor Decrementer exception exists and each attempt to execute an instruction when the Hypervisor Decrementer interrupt is enabled causes an exception (see the Programming Note below), the Hypervisor Decrementer interrupt is not delayed indefinitely.
If LPES=1 and a Direct External exception exists and each attempt to execute an instruction when this interrupt is enabled causes an exception (see the Programming Note below), the Direct External interrupt is not delayed indefinitely.

## Programming Note

An incorrect or malicious operating system could corrupt the first instruction in the interrupt vector location for an instruction-caused interrupt such that the attempt to execute the instruction causes the same exception that caused the interrupt (a looping interrupt; e.g., Trap instruction and Program interrupt). Similarly, the first instruction of the interrupt vector for one instruction-caused interrupt could cause a different instruction-caused interrupt, and the first instruction of the interrupt vector for the second instruction-caused interrupt could cause the first instruction-caused interrupt (e.g., Program interrupt and Floating-Point Unavailable interrupt). Similarly, if the Real Mode Area is virtualized and there is no PTE for the page containing the interrupt vectors, every attempt to execute the first instruction of the OS's Instruction Storage interrupt handler would cause a Hypervisor Instruction Storage interrupt; if the Hypervisor Instruction Storage interrupt handler returns to the OS's Instruction Storage interrupt handler without the relevant PTE having been created, another Hypervisor Instruction Storage interrupt would occur immediately. The looping caused by these and similar cases is terminated by the occurrence of a System Reset or Hypervisor Decrementer interrupt.

### 6.10 Relationship of Event-Based Branches to Interrupts

### 6.10.1 EBB Exception Priority

Event-based branches have a priority lower than that of all interrupts. When an event-based exception is created, the Event-Based Branch facility waits for all possible exceptions that would cause interrupts to be reported. It then generates the event-based branch if no exception that would cause an interrupt exists when the event-based branch is to be generated.

### 6.10.2 EBB Synchronization

When an event-based branch occurs, EBBRR is set to point to an instruction such that all preceding instructions have completed execution, no subsequent instruction has begun execution, and the instruction addressed by EBBRR has not completed execution.

### 6.10.3 EBB Classes

Event-based branches are classified by whether they are directly caused by the execution of an instruction or are caused by some other system exception. Those that are "system-caused" are

- Performance Monitor
- External

The event-based branch caused by execution of an Idmx instruction is "instruction-caused." When execution of an Idmx instruction causes an event-based branch, the following conditions exist.

1. EBBRR addresses the Idmx instruction.
2. An event-based branch is generated such that all instructions preceding the Idmx instruction appear to have completed with respect to the executing thread.
3. The Idmx instruction has not completed.
4. Architecturally, no subsequent instruction has begun execution.

## Chapter 7. Timer Facilities

### 7.1 Overview

The Time Base, Decrementer, Hypervisor Decrementer, Processor Utilization of Resources, and Scaled Processor Utilization of Resources registers provide timing functions for the system. The remainder of this section describes these registers and related facilities.

### 7.2 Time Base (TB)

The Time Base (TB) is a 64-bit register (see Figure 66) containing a 64-bit unsigned integer that is incremented periodically.

| 39 |  |  |  |
| :--- | :--- | :--- | :---: |
|  | TBU40 | TBL |  |
|  | TBU |  |  |
| 0 | 32 |  |  |
| Field | Description |  |  |
| TBU40 | Upper 40 bits of Time Base |  |  |
| TBU | Upper 32 bits of Time Base |  |  |
| TBL | Lower 32 bits of Time Base |  |  |

Figure 66. Time Base
The Time Base is a hypervisor resource; see Chapter 2.

The SPRs TBU40, TBU, and TBL provide access to the fields of the Time Base shown in Figure 66. When a mtspr instruction is executed specifying one of these SPRs, the associated field of the Time Base is altered and the remaining bits of the Time Base are not affected.

See Chapter 6 of Book II for infromation about the update frequency of the Time Base.

The Time Base is implemented such that:

1. Loading a GPR from the Time Base has no effect on the accuracy of the Time Base.
2. Copying the contents of a GPR to the Time Base replaces the contents of the Time Base with the contents of the GPR.

The Power ISA does not specify a relationship between the frequency at which the Time Base is updated and other frequencies, such as the CPU clock or bus clock in a Power ISA system. The Time Base update frequency is not required to be constant. What is required, so that system software can keep time of day and operate interval timers, is one of the following.

■ The system provides an (implementation-dependent) interrupt to software whenever the update frequency of the Time Base changes, and a means to determine what the current update frequency is.

■ The update frequency of the Time Base is under the control of the system software.

Implementations must provide a means for either preventing the Time Base from incrementing or preventing it from being read in problem state $\left(\mathrm{MSR}_{\mathrm{PR}}=1\right)$. If the means is under software control, it must be accessible only in hypervisor state $\left(\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR}=0 \mathrm{~b} 10\right)$. There must be a method for getting all Time Bases in the system to start incrementing with values that are identical or almost identical.

## Programming Note

If software initializes the Time Base on power-on to some reasonable value and the update frequency of the Time Base is constant, the Time Base can be used as a source of values that increase at a constant rate, such as for time stamps in trace entries.

Even if the update frequency is not constant, values read from the Time Base are monotonically increasing (except when the Time Base wraps from $2^{64}-1$ to 0 ). If a trace entry is recorded each time the update frequency changes, the sequence of Time Base values can be post-processed to become actual time values.

Successive readings of the Time Base may return identical values.

If Time Base bits 60:63 are used as part of a random number generator, software must account for the fact that these bits are set to $0 \times 0$ only when bit 59 changes state regardless of whether or not they incremented to 0xF since they were previously set to $0 \times 0$.

See the description of the Time Base in Chapter 6 of Book II for ways to compute time of day in POSIX format from the Time Base.

### 7.2.1 Writing the Time Base

Writing the Time Base is privileged, and can be done only in hypervisor state. Reading the Time Base is not privileged; it is discussed in Chapter 6 of Book II.
It is not possible to write the entire 64-bit Time Base using a single instruction. The mttbl and mttbu extended mnemonics write the lower and upper halves of the Time Base (TBL and TBU), respectively, preserving the other half. These are extended mnemonics for the mtspr instruction; Figure 17.

The Time Base can be written by a sequence such as:

| lwz | Rx, upper \# load 64-bit value for |  |
| :--- | :--- | :--- |
| lwz | Ry,lower \# TB into Rx and Ry |  |
| li | $\mathrm{Rz}, 0$ |  |
| mttbl | Rz | \# set TBL to 0 |
| mttbu | Rx | \# set TBU |
| mttbl | Ry | \# set TBL |

Provided that no interrupts occur while the last three instructions are being executed, loading 0 into TBL prevents the possibility of a carry from TBL to TBU while the Time Base is being initialized.
The preferred method of changing the Time Base utilizes the TBU40 facility. The following code sequence demonstrates the process. Assume the upper 40 bits of Rx contain the desired value upper 40 bits of the Time Base.

| mftb | Ry | \# Read 64-bit Time Base value |
| :---: | :---: | :---: |
| clrldi | Ry, Ry, 40 | \# lower 24 bits of old TB |
| mttbu40 | Rx | \# write upper 40 bits of TB |
| mftb | Rz | \# read TB value again |
| clrldi | Rz, Rz, 40 | \# lower 24 bits of new TB |
| cmpld | Rz, Ry | \# compare new and old lwr 24 |
| bge | done | \# no carry out of low 24 bits |
| addis | Rx, Rx, 0x | 100 |
|  |  | \#increment upper 40 bits |
| mttbu40 | Rx | \# update to adjust for carry |

## Programming Note

The instructions for writing the Time Base are mode-independent. Thus code written to set the Time Base will work correctly in either 64-bit or 32-bit mode.

### 7.3 Virtual Time Base

The Virtual Time Base (VTB) is a 64-bit incrementing counter.


## Figure 67. Virtual Time Base

Virtual Time Base increments at the same rate as the Time Base until its value becomes 0xFFFF_FFFFF_FFFF_FFFF $\left(2^{64}-1\right)$; at the next increment its value becomes $0 \times 0000 \_0000 \_0000 \_0000$. There is no interrupt or other indication when this occurs.

The operation of the Virtual Time Base has the following additional properties.

1. Loading a GPR from the Virtual Time Base has no effect on the accuracy of the Virtual Time Base.
2. Copying the contents of a GPR to the Virtual Time Base replaces the contents of the Virtual Time Base with the contents of the GPR.

## Programming Note

In systems that change the Time Base update frequency for purposes such as power management, the Virtual Time Base input frequency will also change. Software must be aware of this in order to set interval timers.

## Programming Note

In configurations in which the hypervisor allows multiple partitions to time-share a processor, the Virtual Time Base can be managed by the hypervisor such that it appears to each partition as if it counts only during the times that the partition is executing.
In order to do this, the hypervisor saves the value of the Virtual Time Base as part of the program context when removing a partition from the processor, and restores it to its previous value when initiating the partition again on the same or another processor.

### 7.4 Decrementer

| The Decrementer (DEC) is a decrementing counter that provides a mechanism for causing a Decrementer interrupt after a programmable delay.

The Decrementer is driven at the same frequency as I the Time Base.

I


Figure 68. Decrementer
The LPCR is used to enable and disable Large Decrementer mode, as defined below. (See Section 2.2.)

When the Decrementer is not in Large Decrementer mode, it behaves as a 32-bit signed integer and operates as follows.

The Decrementer counts down until its value becomes 0x0000_0000_0000_0000; at the next decrement its value becomes $0 \times 0000$ _0000_FFFF_FFFF. When reading the Decrementer using mfspr, bits 0:31 always read back as 0s.

When the contents of $\mathrm{DEC}_{32}$ change from 0 to 1 , a Decrementer exception will come into existence within a reasonable period of time. When the contents of $\mathrm{DEC}_{32}$ change from 1 to 0 , the existing Decrementer exception, if any, will cease to exist within a reasonable period of time, but not later than the completion of the next context synchronizing instruction or event.

The preceding paragraph applies regardless of whether the change in the contents of $\mathrm{DEC}_{32}$ is the result of decrementation of the Decrementer by the hardware or of modification of the Decrementer caused by execution of an mtspr instruction.
When the Decrementer is in Large Decrementer mode, it behaves as a d-bit decrementing counter which is sign-extended to 64 bits. The value of $d$ is implementa-
tion dependent but at least 32. When the Decrementer is written, bits 0:63-d are ignored by the hardware.

## Programming Note

In Large Decrementer mode, the maximum positive value supported by the Decrementer is $2^{\mathrm{d}-1}-1$, represented with bits 0:64-d containing 0's and bits $65-\mathrm{d}: 63$ containing 1's. The minimum value supported by the Decrementer is $-2^{d-1}$, represented as $0 x F F F F \_F F F F \_F F F F \_F F F F$.

When in Large Decrementer mode, the Decrementer operates as follows.

The binary value of the Decrementer counts down until its value becomes 0x0000_0000_0000_0000; at the next decrement its value becomes the minimum value supported, which is represented as 0xFFFF_FFFF_FFFF_FFFF.

When the contents of the $\mathrm{DEC}_{0}$ change from 0 to 1, a Decrementer exception will come into existence within a reasonable period of time. When the contents of $\mathrm{DEC}_{0}$ change from 1 to 0 , the existing Decrementer exception, if any, will cease to exist within a reasonable period of time, but not later than the completion of the next context synchronizing instruction or event.
The preceding paragraph applies regardless of whether the change in the contents of $\mathrm{DEC}_{0}$ is the result of decrementation of the Decrementer by the hardware or of modification of the Decrementer caused by execution of an mtspr instruction.

The operation of the Decrementer has the following additional properties.

1. Loading a GPR from the Decrementer has no effect on the accuracy of the Time Base.
2. Copying the contents of a GPR to the Decrementer replaces the contents of the Decrementer with the contents of the GPR.

## Programming Note

In systems that change the Time Base update frequency for purposes such as power management, the Decrementer input frequency will also change. Software must be aware of this in order to set interval timers.
If Decrementer bits 60:63 are used as part of a random number generator, software must account for the fact that these bits are set to 0xF only when bit 59 changes state regardless of whether or not they decremented to $0 \times 0$ since they were previously set to $0 x F$.

### 7.4.1 Writing and Reading the Decrementer

The contents of the Decrementer can be read or written using the mfspr and mtspr instructions, both of which are privileged when they refer to the Decrementer. Using an extended mnemonic (Figure 17), the Decrementer can be written from GPR Rx using:

```
mtdec Rx
```

The Decrementer can be read into GPR Rx using:

```
mfdec Rx
```

Copying the Decrementer to a GPR has no effect on the Decrementer contents or on the interrupt mechanism.

### 7.5 Hypervisor Decrementer

The Hypervisor Decrementer is a h-bit decrementing counter that is sign-extended to 64 bits. The value of $h$ is implementation dependent, however the number of bits supported by the Hypervisor Decrementer must be greater than or equal to the number of bits supported by the Decrementer. When the Decrementer is written, bits 0:63-h are ignored by the hardware.

## _ Programming Note <br> The maximum positive value supported by the Hypervisor Decrementer is $2^{h-1}-1$, represented with bits 0:64-h containing 0's and bits 65-h:63 containing 1's. The minimum value supported by the Hypervisor Decrementer is $-2^{h-1}$, represented as $0 x F F F F$ _FFFF_FFFF_FFFF.

The binary value of the Hypervisor Decrementer counts down until its value becomes 0x0000_0000_0000_0000; at the next decrement its value becomes the minimum value supported, which is represented as 0xFFFF_FFFF_FFFF_FFFF.

When the contents of $\mathrm{HDEC}_{0}$ change from 0 to 1 and the thread is not in a power-saving mode, a Hypervisor Decrementer exception will come into existence within a reasonable period of time. When a Hypervisor Decrementer interrupt occurs, the existing Hypervisor Decrementer exception will cease to exist within a reasonable period of time, but not later than the completion of the next context synchronizing instruction or event. Even if multiple $\mathrm{HDEC}_{0}$ change transitions from 0 to 1 occur before a Hypervisor Decrementer interrupt occurs, at most one Hypervisor Decrementer exception exists.

The preceding paragraph applies regardless of whether the change in the contents of $\mathrm{HDEC}_{0}$ is the result of decrementation of the Hypervisor Decrementer by the hardware or of modification of the Hypervisor Decrementer caused by execution of an mtspr instruction.

The operation of the Hypervisor Decrementer has the following additional properties.

1. Loading a GPR from the Hypervisor Decrementer has no effect on the accuracy of the Hypervisor Decrementer.
2. Copying the contents of a GPR to the Hypervisor Decrementer replaces the contents of the Hypervisor Decrementer with the contents of the GPR.

## Programming Note

In systems that change the Time Base update frequency for purposes such as power management, the Hypervisor Decrementer update frequency will also change. Software must be aware of this in order to set interval timers.
If Hypervisor Decrementer bits 60:63 are used as part of a random number generator, software must account for the fact that these bits are set to 0xF only when bit 59 changes state regardless of whether or not they decremented to $0 \times 0$ since they were previously set to 0xF.

## Programming Note

A Hypervisor Decrementer exception is not created if the thread is in a power-saving mode when $\mathrm{HDEC}_{0}$ changes from 0 to 1 because having a Hypervisor Decrementer interrupt occur almost immediately after exiting the power-saving mode in this case is deemed unnecessary. The hypervisor already has control, and if a timed exit from the power-saving mode is necessary and possible, the hypervisor can use the Decrementer to exit the power-saving mode at the appropriate time. For some power-saving levels, the state of the Hypervisor Decrementer and Decrementer is not necessarily maintained and updated.

### 7.6 Processor Utilization of Resources Register (PURR)

The Processor Utilization of Resources Register (PURR) is a 64-bit counter, the contents of which provide an estimate of the resources used by the thread. The contents of the PURR are treated as a 64-bit unsigned integer.


Figure 69. Processor Utilization of Resources Register

The PURR is a hypervisor resource; see Chapter 2.
The contents of the PURR increase monotonically, unless altered by software, until the sum of the contents
plus the amount by which it is to be increased exceed OxFFFF_FFFF_FFFF_FFFF $\left(2^{64}-1\right)$ at which point the contents are replaced by that sum modulo $2^{64}$. There is no interrupt or other indication when this occurs.

The rate at which the value represented by the contents of the PURR increases is an estimate of the portion of resources used by the thread per unit time with respect to other threads that share those resources monitored by the PURR. When the thread is idle, the rate at which the PURR value increases is implementation dependent.

Let the difference between the value represented by the contents of the Time Base at times $T_{a}$ and $T_{b}$ be $\mathrm{T}_{\mathrm{ab}}$. Let the difference between the value represented by the contents of the PURR at time $T_{a}$ and $T_{b}$ be the value $P_{a b}$. The ratio of $P_{a b} / T_{a b}$ is an estimate of the percentage of shared resources used by the thread during the interval $T_{a b}$. For the set $\{S\}$ of threads that share the resources monitored by the PURR, the sum of the usage estimates for all the threads in the set is 1.0 .

The definition of the set of threads $S$, the shared resources corresponding to the set $S$, and specifics of the algorithm for incrementing the PURR are imple-mentation-specific.

The PURR is implemented such that:

1. Loading a GPR from the PURR has no effect on the accuracy of the PURR.
2. Copying the contents of a GPR to the PURR replaces the contents of the PURR with the contents of the GPR.

## Programming Note

Estimates computed as described above may be useful for purposes related to resource utilization, including utilization-based system management and planning.
Because the rate at which the PURR accumulates resource usage estimates is dependent on the frequency at which the Time Base is incremented, and the frequency of the oscillator that drives instruction execution may vary independently from that of the Time Base, the interpretation of the contents of the PURR may be inaccurate as a measurement of capacity consumption for accounting purposes. The SPURR should be used for accounting purposes.

### 7.7 Scaled Processor Utilization of Resources Register (SPURR)

The Scaled Processor Utilization of Resources Register (SPURR) is a 64-bit counter, the contents of which provide an estimate of the resources used by the
thread. The contents of the SPURR are treated as a 64-bit unsigned integer.


The SPURR is a hypervisor resource; see Section 2.6. The contents of the SPURR increase monotonically, unless altered by software, until the sum of the contents plus the amount by which it is to be increased exceed OxFFFF_FFFF_FFFF_FFFF $\left(2^{64}-1\right)$ at which point the contents are replaced by that sum modulo $2^{64}$. There is no interrupt or other indication when this occurs.
The rate at which the value represented by the contents of the SPURR increases is an estimate of the portion of resources used by the thread with respect to other threads that share those resources monitored by the SPURR, and relative to the computational capacity provided by those resources. The computational capacity provided by the shared resources may vary as a function of the frequency of the oscillator which drives the resources or as a result of deliberate delays in processing that are created to reduce power consumption. When the thread is idle, the rate at which the SPURR value increases is implementation dependent.
Let the difference between the value represented by the contents of the Time Base at times $T_{a}$ and $T_{b}$ be $\mathrm{T}_{\mathrm{ab}}$. Let the ratio of the effective and nominal frequencies of the oscillator driving instruction execution $f_{e} / f_{n}$ be $f_{r}$. Let the ratio of delay cycles created by power reduction circuitry and total cycles $c_{d} / c_{t}$ be $c_{r}$. Let the difference between the value represented by the contents of the SPURR at time $T_{a}$ and $T_{b}$ be the value $S_{a b}$. The ratio of $S_{a b} /\left(T_{a b} \times f_{r} \times\left(1-c_{r}\right)\right)$ is an estimate of the percentage of shared resource capacity used by the thread during the interval $T_{a b}$. For the set $\{S\}$ of threads that share the resources monitored by the SPURR, the sum of the usage estimates for all the threads in the set is 1.0 .

The definition of the set of threads $S$, the shared resources corresponding to the set S , and specifics of the algorithm for incrementing the SPURR are imple-mentation-specific.
The SPURR is implemented such that:

1. Loading a GPR from the SPURR has no effect on the accuracy of the SPURR.
2. Copying the contents of a GPR to the SPURR replaces the contents of the SPURR with the contents of the GPR.

## Version 3.0

## Programming Note

Estimates computed as described above may be useful for purposes of resource use accounting, program dispatching, etc.

### 7.8 Instruction Counter

The Instruction Counter (IC) is a 64-bit incrementing counter that counts the number of instructions that the thread has completed (according to the sequential execution model; see Section 2.2 of Book I).


Figure 71. Instruction Counter

# Chapter 8. Debug Facilities 

### 8.1 Overview

Implementations provide debug facilities to enable hardware and software debug functions, such as control flow tracing, data address watchpoints, and program single-stepping. The debug facilities described in this section consist of the Come-From Address Register (see Section 8.2), Completed Instruction Address Breakpoint Register (see Section 8.3), and the Data Address Watchpoint Register (DAWRn) and Data Address Watchpoint Register Extension (DAWRXn) (see Section 8.4). The interrupt associated with the Data Address Breakpoint registers is described in Section 6.5.3. The interrupt associated with the Completed Instruction Address Breakpoint Register is described in Section 6.5.15. The Trace facility, which can be used for single-stepping as well as for control flow tracing, is described in Section 6.5.15.

The mfspr and mtspr instructions (see Section 4.4.5) provide access to the registers of the debug facilities.

In addition to the facilities mentioned above, implementations typically provide debug facilities, modes, and access mechanisms that are implementation-specific. For example, implementations typically provide facilities for instruction address tracing, and also access to certain debug facilities via a dedicated interface such as the IEEE 1149.1 Test Access Port (JTAG).

### 8.2 Come-From Address Register

The Come-From Address Register (CFAR) is a 64-bit | register. When an rfebb, rfid, or rfscv instruction is executed, the register is set to the effective address of the instruction. When a Branch instruction is executed and the branch is taken, the register is set to the effective address of an instruction in the instruction cache block containing the Branch instruction, except that if the Branch instruction is a B-form Branch (i.e., bc, bca, $\boldsymbol{b c l}$, or bcla) for which the target address is in the instruction cache block containing the Branch instruction or is in the previous or next cache block, the register is not necessarily set. For Branch instructions, the
setting need not occur until a subsequent context synchronizing operation has occurred.


## Figure 72. Come-From Address Register

The contents of the CFAR can be read and written using the mfspr and mtspr instructions. Acccess to the CFAR is privileged.

## Programming Note

This register can be used for purposes of debugging software. For example, often a software bug results in the program executing a portion of the code that it should not have reached or causing an unexpected interrupt. In the former case, a breakpoint can be placed in the portion of the code that was erroneously reached and the program reexecuted. In either case, the interrupt handler can save the contents of the CFAR (before executing the first instruction that would modify the register), and then make the saved contents available for a debugger to use in determining the control flow path by which the exception was reached.
In order to preserve the CFAR's contents for each partition and to prevent it from being used to implement a "covert channel" between partitions, the hypervisor should initialize/save/restore the CFAR when switching partitions on a given thread.

### 8.3 Completed Instruction Address Breakpoint

The Completed Instruction Address Breakpoint mechanism provides a means of detecting an instruction completion at a specific instruction address. The address comparison is done on an effective address (EA).

The Completed Instruction Address Breakpoint mechanism is controlled by the Completed Instruction

Address Breakpoint Register (CIABR), shown in Figure 74.

|  | CIEA |
| :---: | :---: |
| 0 | PRIV |

$\begin{array}{ll}\text { Bit(s) } & \text { Name } \\ 0: 61 & \text { CIEA }\end{array}$

## Description

0:61 CIEA Completed Instruction Effective Address
62:63 PRIV
Privilege
00: Disable matching
01: Match in problem state
10: Match in privileged (non-hypervisor) state
11: Match in hypervisor state
Figure 73. Completed Instruction Address Breakpoint Register

A Completed Instruction Address Breakpoint match occurs upon instruction completion if all of the following conditions are satisfied.

■ the completed instruction address is equal to CIEA $_{0: 61}$ II $0 b 00$.

- the thread run level matches that specified in RLM.

In 32-bit mode the high-order 32 bits of the EA are treated as zeros for the purpose of detecting a match.
A Completed Instruction Address Breakpoint match causes a Trace exception provided that no higher priority interrupt occurs from the completion of the instruction (see Section 6.5.15).

### 8.4 Data Address Watchpoint

The Data Address Watchpoint mechanism provides a means of detecting load and store accesses to a range of addresses starting at a designated doubleword. The address comparison is done on an effective address (EA).

## - Programming Note

The Data Address Watchpoint mechanism employs a simple EA compare. It makes no attempt to take the radix table translation quadrants (keyed off $E A_{0: 1}$ ) into account to enable a single setting to work in all privilege levels.

The Data Address Watchpoint mechanism is controlled by a single set of SPRs, numbered with $\mathrm{n}=0$ : the Data Address Watchpoint Register (DAWRn), shown in

Figure 74, and the Data Address Watchpoint Register Extension (DAWRXn), shown in Figure 75.


Bit(s) Name Description
0:60 DEAW Data Effective Address Watchpoint
Figure 74. Data Address Watchpoint Register

| //I | MRD | /// | HRAMMC | DW | DR | WT | WTI | PRIVM |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 32 | 48 | 54 | 56 | 57 | 58 | 59 | 60 | 61 |


| Bit(s) | Name | Description |
| :---: | :---: | :---: |
| 48:53 | MRD | Match Range in Doublewords biased by $-1 .(0 b 000000=1$ DW Ob111111 = 64 DW) |
| 56 | HRAMMC | Hypervisor Real Addressing Mod Match Control <br> 0 : DEAW ${ }_{0}$ and $E A_{0}$ are used during matching in hypervisor real addressing mode <br> 1: DEAW $_{0}$ and EA ${ }_{0}$ are ignored during matching in hypervisor real addressing mode |
| 57 | DW | Data Write |
| 58 | DR | Data Read |
| 59 | WT | Watchpoint Translation |
| 60 | WTI | Watchpoint Translation Ignore |
| 61:63 | PRIVM | Privilege Mask |
| 61 | HYP | Hypervisor state |
| 62 | PNH | Privileged but Non-Hypervisor state |
| 63 | PRO | Problem state |

All other fields are reserved.
Figure 75. Data Address Watchpoint Register Extension

The supported PRIVM values are 0b000, 0b001, 0b010, Ob011, Ob100, and 0b111. If the PRIVM field does not contain one of the supported values, then whether a match occurs for a given storage access is undefined. Elsewhere in this section it is assumed that the PRIVM field contains one of the supported values.

## Programming Note

PRIVM value Ob000 causes matches not to occur regardless of the contents of other DAWRn and DAWRXn fields. PRIVM values Ob101 and Ob110 are not supported because a storage location that is shared between the hypervisor and non-hypervisor software is unlikely to be accessed using the same EA by both the hypervisor and the non-hypervisor software. (PRIVM value 0b111 is supported primarily for reasons of software compatibility with respect to emulation of the DABR facility as described in a subsequent Programming Note.)

A Data Address Watchpoint match occurs for a Load or Store instruction if, for any byte accessed, all of the following conditions are satisfied.

■ the access is

- a quadword access and located in the range $\left(\right.$ DEAW $\left._{0: 59} \| \operatorname{ObO}\right) \leq\left(E_{0: 59} \| \mathrm{ObO}^{2}\right) \leq$ ((DEAW $0: 59$ II ObO) $+\left({ }^{55} 0\right.$ II $M R D_{0: 4}$ Il $\left.0 b 0\right)$ ) such that $\left(E A_{0: 60}\right.$ AND $\left.\left({ }^{55} 1 \|{ }^{6} 0\right)\right)=$ (DEAW ${ }_{0: 60}$ AND $\left({ }^{55} 1\right.$ II $\left.{ }^{6} 0\right)$ ).
- not a quadword access and located in the range $\mathrm{DEAW}_{0: 60} \leq \mathrm{EA}_{0: 60} \leq$ $\left(\right.$ DEAW $_{0: 60}+{ }^{55}{ }^{55}$ II $\left.M R D_{0: 5}\right)$ ) such that $\left(E A_{0: 60}\right.$ AND $\left.\left({ }^{55} 1 \|{ }^{6} 0\right)\right)=$ (DEAW ${ }_{0: 60}$ AND $\left({ }^{55} 1\right.$ || $\left.{ }^{6} 0\right)$ ).
- ( $\left.\mathrm{MSR}_{\mathrm{DR}}=\mathrm{DAWRX}_{\mathrm{WT}}\right) \mid$ DAWRXn $\mathrm{n}_{\mathrm{WTI}}$
- the thread is in
- hypervisor state and DAWRXn $n_{H Y P}=1$, or
- privileged but non-hypervisor state and DAWRXn $_{\text {PNH }}=1$, or
- problem state and DAWRXn $n_{P R}=1$

■ the instruction is a Store and DAWRXn $n_{D W}=1$, or the instruction is a Load and DAWRXn $n_{D R}=1$.

In 32-bit mode the high-order 32 bits of the EA are treated as zeros for the purpose of detecting a match.

If the above conditions are satisfied, it is undefined whether a match occurs in the following cases.

■ The instruction is Store Conditional but the store is not performed

- The instruction is dcbz. (For the purpose of determining whether a match occurs, dcbz is treated as a Store.)
The Cache Management instructions other than dcbz never cause a match.

A Data Address Watchpoint match causes a Data Storage exception or a Hypervisor Data Storage exception (see Section 6.5.3, "Data Storage Interrupt" on page 1065 and Section 6.5.16, "Hypervisor Data Storage Interrupt" on page 1075). If a match occurs, some or all of the bytes of the storage operand may have been accessed; however, if a Store instruction causes the match, the storage operand is not modified if the instruction is one of the following:

- any Store instruction that causes an atomic access


## Programming Note

The Data Address Watchpoint mechanism does not apply to instruction fetches.

## Programming Note

Implementations that comply with versions of the architecture that precede Version 2.02 do not provide the DABRX (now replaced by DAWRXn). Forward compatibility for software that was written for such implementations (and uses the Data Address Breakpoint facility) can be obtained by setting $\mathrm{DAWRXn}_{60: 63}$ to 0b0111.

Version 3.0

1100 Power ISA ${ }^{\text {TM }}$ - Book III

# Chapter 9. Performance Monitor Facility 

### 9.1 Overview

The Performance Monitor facility provides a means of collecting information about program and system performance.

### 9.2 Performance Monitor Operation

The Performance Monitor facility includes the following features.

- an MSR bit
- PMM (Performance Monitor Mark), which can be used to select one or more programs for monitoring
- registers
- PMC1 - PMC6 (Performance Monitor Counters 1-6), which count events
- MMCR0, MMCR1, MMCR2, and MMCRA (Monitor Mode Control Registers 0, 1, 2, and A), which control the Performance Monitor facility
- SIAR, SDAR, and SIER (Sampled Instruction Address Register, Sampled Data Address Register, and Sampled Instruction Event Register), which contain the address of the "sampled instruction" and of the "sampled data," and additional information about the "sampled instruction" (see Section 9.4.8 - Section 9.4.10).

■ the Performance Monitor interrupt and Performance Monitor event-based branch, which can be caused by monitored conditions and events.

Many aspects of the operation of the Performance Monitor are summarized by the following hierarchy, which is described starting at the lowest level.

■ A "counter negative condition" exists when the value in a PMC is negative (i.e., when bit 0 of the PMC is 1). A "Time Base transition event" occurs
when a selected bit of the Time Base changes from 0 to 1 (the bit is selected by a field in MMCRO). The term "condition or event" is used as an abbreviation for "counter negative condition or Time Base transition event". A condition or event can be caused implicitly by the hardware (e.g., incrementing a PMC) or explicitly by software (mtspr).

- A condition or event is enabled if the corresponding "Enable" bit (i.e., PMC1CE, PMCjCE, or TBEE) in MMCRO is 1 . The occurrence of an enabled condition or event can have side effects within the Performance Monitor, such as causing the PMCs to cease counting.
- An enabled condition or event causes a Performance Monitor alert if Performance Monitor alerts are enabled by the corresponding "Enable" bit in MMCRO. Another cause of a Performance Monitor alert is the threshold event counter reaching its maximum value (see Section 9.4.3). A single Performance Monitor alert may reflect multiple enabled conditions and events.

■ When a Performance Monitor alert occurs, MMCR $_{\text {PMAO }}$ is set to 1 and the writing of BHRB entries, if in process, is suspended.

When the contents of MMCRO $0_{\text {PMAO }}$ change from 0 to 1, a Performance Monitor exception will come into existence within a reasonable period of time. When the contents of MMCRO PMAO change from 1 to 0 , the existing Performance Monitor exception, if any, will cease to exist within a reasonable period of time, but not later than the completion of the next context synchronizing instruction or event.

- A Performance Monitor exception causes one of the following.
- If $M S R_{E E}=1, M M C R 0_{E B E}=0$, and either $\operatorname{HFSCR}_{\text {PM }}=1$ or the thread is in hypervisor state, an interrupt occurs.
- If $M S R_{P R}=1, M M C R 0^{E B E}=1$, a Performance Monitor event-based exception occurs if BESCR $_{\text {PME }}=1$, provided that event-based exceptions are enabled by FSCR $_{\text {EBB }}$ and HFSCR $_{\text {EbB }}$. When a Performance Monitor
event-based exception occurs, an event-based branch is generated if BES$\mathrm{CR}_{\mathrm{GE}}=1$.


## Programming Note

The Performance Monitor can be effectively disabled (i.e., put into a state in which Performance Monitor SPRs are not altered and Performance Monitor exceptions do not occur) by setting MMCR0 to 0x0000_0000_8000_0000.

The Performance Monitor also controls when BHRB entries are written, the instruction filters that are used when writing BHRB entries, and the availability of the BHRB in problem state. It also controls whether Performance Monitor exceptions cause Performance Monitor event-based exceptions or Performance Monitor interrupts. See Section 9.4.4.

### 9.3 No-op Instructions Reserved for the Performance Monitor

The following forms of the and $x, x, x$ instruction are reserved for exclusive use by the Performance Monitor.

- and $\mathrm{x}, \mathrm{x}, \mathrm{x}$, where $\mathrm{x}=0,1$.


## - Programming Note

An example usage of a probe no-op by the Performance Monitor is to measure branch prediction effectiveness. In order to do this, one of probe no-ops is inserted in various sections of the code in which branch prediction efficiency is being studied. The Performance Monitor registers are then set up as follows.

## MMCRA:

ES=010 (only probe no-ops eligible for sampling)
SM=00 (all eligible instructions)
$\mathrm{SE}=1$ (enable random sampling).
Other fields in MMCRA are set as desired.

## MMCR1:

PMC1SEL=E0 (count PMC1 on dispatch)
PMC4SEL=E0 (count PMC4 on completion)
Other counters initialized as desired.
MMCR2: Initialize as desired.

## MMCRO:

FC is set to 0 to stop freezing the counters
PMAE is set to 1 to enable PMU alerts.
Other fields in MMCRO are set as desired.
Subsequently, when a PMU alert occurs, PMCs 1 and 4 can be read. The difference between the two counter values provides an indication of branch prediction effectiveness in the areas of the code in which the probe no-op was inserted.

### 9.4 Performance Monitor Facility Registers

The Performance Monitor registers count events, control the operation of the Performance Monitor, and provide associated information.

The elapsed time between the execution of an instruction and the time at which events due to that instruction have been reflected in Performance Monitor registers is not defined. No means are provided by which software can ensure that all events due to preceding instructions have been reflected in Performance Monitor registers. Similarly, if the events being monitored may be caused by operations that are performed out-of-order, no means are provided by which software can prevent such events due to subsequent instructions from being reflected in Performance Monitor registers. Thus the contents obtained by reading a Performance Monitor register may not be precise: it may fail to reflect some events due to instructions that precede the mfspr and may reflect some events due to instructions that follow the mfspr. This lack of precision applies regardless of whether the state of the thread is such that the register is subject to change by the hardware at the time the $\boldsymbol{m f s p r}$ is executed. Similarly, if an mtspr instruction is executed that changes the contents of the Time Base, the change is not guaranteed to have taken effect with respect to causing Time Base transition events until after a subsequent context synchronizing instruction has been executed.

If an mtspr instruction is executed that changes the value of a Performance Monitor register other than SIAR, SDAR, and SIER, the change is not guaranteed to have taken effect until after a subsequent context synchronizing instruction has been executed (see Chapter 11. "Synchronization Requirements for Context Alterations" on page 1127).

## - Programming Note

Depending on the events being monitored, the contents of Performance Monitor registers may be affected by aspects of the runtime environment (e.g., cache contents) that are not directly attributable to the programs being monitored.

### 9.4.1 Performance Monitor SPR Numbers

The Performance Monitor registers have two sets of SPR numbers, one set that is non-privileged and another set that is privileged.
For the purpose of explanation elsewhere in the architecture, the non-privileged registers are divided into two groups as defined below.

■ A: The non-privileged read/write Performance Monitor registers (i.e., the PMCs, MMCRO, MMCR2, and MMCRA at SPR numbers 771-776, 779, 769, and 770, respectively)

- B: The non-privileged read-only Performance Monitor registers (i.e., SIER, SIAR, SDAR, and MMCR1 at SPR numbers 768, 780, 781, and 782, respectively).

The SPRs in group B are treated as not implemented registers for write (mtspr) operations. See the mtspr instruction description in Section 4.4.5 for additional information.

When the PCR makes a register in either group A or B unavailable in problem state, that SPR is not included in group $A$ or $B$.

## Programming Note

Older versions of Performance Monitor facilities used diffefrent sets of SPR numbers from those shown in Section 4.4.5. (All 32-bit PowerPC implementations used a different set.

### 9.4.2 Performance Monitor Counters

The six Performance Monitor Counters, PMC1 through PMC6, are 32-bit registers that count events.

| PMC1 |
| :---: |
| PMC2 |
| PMC3 |
| PMC4 |
| PMC5 |
| PMC6 |
| 32 |

Figure 76. Performance Monitor Counter registers
PMC1 - PMC4 are referred to as "programmable" counters since the events that can be counted can be specified by the program. The events that are counted by each counter are specified in MMCR1.

PMC5 and PMC6 are not programmable and can be specified as being part of the Performance Monitor Facility or not part of it. PMC5 counts instructions completed, and PMC6 counts cycles. The PMCC field in MMCRO controls whether or not PMCs 5-6 are part of the Performance Monitor Facility, and the result of accessing these counters when they are not part of the Performance Monitor Facility.

## Programming Note

PMC5 and PMC6 are defined to facilitate calculating basic performance metrics such as cycles per instruction (CPI).

## Programming Note

Software can use a PMC to "pace" the collection of Performance Monitor data. For example, if it is desired to collect event counts every $n$ cycles, software can specify that a particular PMC count cycles, and set that PMC to 0x8000_0000-n. The events of interest would be counted in other PMCs. The counter negative condition that will occur after n cycles can, with the appropriate setting of MMCR bits, cause counter values to become frozen, cause a Performance Monitor exception to occur, etc.

### 9.4.2.1 Event Counting and Sampling

The PMCs are enabled to count unless they are "frozen" by one or more of the "freeze counters" fields in MMCRO or MMCR2.

Each of PMC's 1-4 can be configured, using MMCR1, to count "continuous" events (events that can occur at any time), or to count "randomly sampled" events (or "sampled" events) that are associated with the execution of randomly sampled instructions.
Continuous events always cause the counters to count (unless counters are frozen). These events are specified for each counter by using encodes F0-FF in the PMCn Selector fields in MMCR1.

Randomly sampled events can cause the counters to count only when random sampling has been enabled by setting $\mathrm{MMCRO}_{\mathrm{SE}}=1$. The types of instructions that are sampled are specified in MMCRA SM and MMCRA $_{\text {Es }}$. Randomly sampled events are specified for each counter by using encodes E0-EF in the PMCn Selector fields in MMCR1.

## Programming Note

A typical sequence of operations that enables use the PMCs is as follows.

- Freeze the counters by setting $M M C R 0_{\mathrm{FC}}=1$.
- Set control fields in MMCR0 and MMCR2 that control counting in various privilege states and other modes, and that enable counter negative conditions.
- Initialize the events to be counted by PMCs 1-4 using the PMCn Selector fields in MMCR1.
- Specify the BHRB filtering mode, threshold event Counter events, and whether or not random sampling is enabled in the corresponding fields in MMCRA.
- Initialize the PMCs to the values desired. For example, in order to configure a counter to cause a counter negative condition after $n$ counts, that counter would be initialized to $2^{32}$-n.
- Set $\mathrm{MMCRO}_{\mathrm{FC}}$ to 0 to disable freezing the counters, and set MMCRO $0_{\text {PMAE }}$ to 1 if a Performance Monitor alert (and the corresponding Performance Monitor interrupt) is desired when an enabled condition or event occurs. (See Section 9.2 for the definition of enabled condition or event.)
When the Performance Monitor alert occurs, the program would typically read the values of the counters as well as the contents of SIAR, SDAR, SIER as needed in order to extract the information that was being monitored.

See Sections 9.4.4-9.4.10 for information regarding MMCRs, SIAR, SDAR, and SIER, and some additional usage examples.

### 9.4.3 Threshold Event Counter

The threshold event counter and associated controls are in MMCRA (see Section 9.4.7). When Performance Monitor alerts are enabled (MMCRO $0_{\text {PMAE }}=1$ ), this counter begins incrementing from value 0 upon each occurrence of the event specified in the Threshold Event Counter Event (TECE) field after the event specified by the Threshold Start Event (TS) field occurs. The counter stops incrementing when the event specified in the Threshold End Event (TE) field occurs. The counter subsequently freezes until the event specified in the TS field is again recognized, at which point it restarts incrementing from value 0 as explained above. If the counter reaches its maximum value or a Performance Monitor alert occurs, incrementing stops. After the Performance Monitor alert occurs, the contents of the threshold event counter are not altered by the hardware until software sets MMCRO $0_{\text {PMAE }}$ to 1 .

## Programming Note

Because hardware can modify the contents of the threshold event counter when random sampling is enabled (MMCRA SE $=1$ ) and MMCR0 ${ }_{\text {PMAE }}=1$ at any time, any value written to the threshold event counter under this condition may be immediately overwritten by hardware.

The threshold event counter value is represented as a 3 -bit integral power of 4 , multiplied by a 7 -bit integer. The exponent is contained in MMCRA ${ }_{\text {TECX }}$, and the multiplier is contained in MMCRA TECM. . For a given counter exponent, $e$, and multiplier, $m$, the number represented is as follows:

$$
\mathrm{N}=4^{\mathrm{e}} \times \mathrm{m}
$$

This counter format allows the counter to represent a range of 0 through approximately 2 million counts with many fewer bits than would be required by a binary counter.

To represent a given counter value, hardware uses as e the smallest 3 -bit integer for which a 7-bit integer exists such that the given counter value can be expressed using this format.

## - Programming Note <br> Software can obtain the number N from the contents of the threshold event counter by shifting the multiplier left twice times the value contained in the exponent.

The value in the counter is the exact number of events that occur for values from 0 through the maximum multiplier value (127), within 4 events of the exact value for values from 128-508 (or $127 \times 4$ ), within 16 events of the exact value for values from 512-2032 (or $127 \times 4^{2}$ ), and so on. This represents an event count accuracy of approximately $3 \%$, which is expected to be sufficient for most situations in which a count of events between a start and end event is required.

## Programming Note

When using the threshold event counter, software typically specifies a "threshold counter exceeded n" event in MMCR1. This enables a PMC to count the number of times the counter exceeded a specified threshold value during the time Performance Monitor alerts were enabled.

### 9.4.4 Monitor Mode Control Register 0

Monitor Mode Control Register 0 (MMCRO) is a 64-bit register as shown below.


Figure 77. Monitor Mode Control Register 0
MMCR0 is used to control multiple functions of the Performance Monitor. Some fields of MMCR0 are altered by the hardware when various events occur.

The following notation is used in the definitions below. "PMCs" refers to PMCs $1-\mathrm{n}$ and "PMCj" refers to PMCj, where $2 \leq \mathrm{j} \leq \mathrm{n}$. $\mathrm{n}=4$ when MMCRO ${ }_{\text {PMCC }}=0 \mathrm{~b} 11$ and $n=6$ otherwise.

When MMCR0 PMCC is set to $0 b 10$ or $0 b 11$, providing problem state programs read/write access to MMCRO, only FC, PMAE, PMAO can be accessed. All other bits are not changed when mtspr is executed in problem state, and all other bits return 0s when mfspr is executed in problem state.

## Programming Note

When PMCC=0b10 or Ob11, problem state programs have write access to MMCRO in order to enable event-based branch routines to reset the FC bit after it has been set to 1 as a result of an enabled condition or event (FCECE=1). During event processing, the event-based branch handler would write the desired initial values to the PMCs and reset the FC bit to 0. PMAO and PMAE can also be set to their appropriate values during the same write operation before returning.

The bit definitions of MMCRO are as follows.
Bit(s) Description
0:31 Reserved
32 Freeze Counters (FC)
0 The PMCs are incremented (if permitted by other MMCR bits).
1 The PMCs are not incremented.
The hardware sets this bit to 1 when an enabled condition or event occurs and $M_{M C R} 0_{\text {FCECE }}=1$.

33
Freeze Counters and BHRB in Privileged State (FCS)

0 The PMCs are incremented (if permitted by other MMCR bits), and entries are written into the BHRB (if permitted by the BHRB Instruction Filtering Mode field in MMCRA).

1 The PMCs are not incremented, and entries are not written into the BHRB, if $\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR}^{=}=0 \mathrm{~b} 00$.

## Conditionally Freeze Counters and BHRB in Problem State (FCP)

If the value of bit 51 (FCPC) is 0 , this field has the following meaning.

0 The PMCs are incremented (if permitted by other MMCR bits) and entries are written into the BHRB (if permitted by the BHRB Instruction Filtering Mode field in MMCRA).
1 The PMCs are not incremented, and entries are not written into the BHRB, if $M_{\text {PR }}=1$.
If the value of bit 51 (FCPC) is 1 , this field has the following meaning.
0 The PMCs are not incremented, and entries are not written into the BHRB, if $\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR}^{2}=0 \mathrm{~b} 01$.
1 The PMCs are not incremented, and entries are not written into the BHRB, if $\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR}=0 \mathrm{~b} 11$.

## Programming Note

In order to freeze counters in problem state regardless of $\mathrm{MSR}_{\mathrm{HV}}, \mathrm{MMCRO}_{\mathrm{FCPC}}$ must be set to 0 and $\mathrm{MMCRO}_{\mathrm{FCP}}$ must be set to 1 .

Freeze Counters while Mark = 1 (FCM1)
0 The PMCs are incremented (if permitted by other MMCR bits).
1 The PMCs are not incremented if $\mathrm{MSR}_{\mathrm{PMM}}=1$.

## Freeze Counters while Mark = $\mathbf{0}$ (FCMO)

0 The PMCs are incremented (if permitted by other MMCR bits).
1 The PMCs are not incremented if $M_{\text {PRM }}=0$.

Performance Monitor Alert Enable (PMAE)
0 Performance Monitor alerts are disabled and BHRB entries are not written.
1 Performance Monitor alerts are enabled, and BHRB entries are written (if enabled by other bits) until a Performance Monitor alert occurs, at which time:

- MMCRO $0_{\text {PMAE }}$ is set to 0
- MMCRO $0_{\text {PMAO }}$ is set to 1

39:40 Time Base Selector (TBSEL)
This field selects the Time Base bit that can cause a Time Base transition event (the event occurs when the selected bit changes from 0 to 1).
00 Time Base bit 63 is selected.
01 Time Base bit 55 is selected.
10 Time Base bit 51 is selected.
11 Time Base bit 47 is selected. condition or event occurs when $\mathrm{MMCRO}_{\text {TRIGGER }}=0$, at which time:

- MMCRO $0_{F C}$ is set to 1

If the enabled condition or event occurs when MMCR $_{\text {TRIGGER }}=1$, the FCECE bit is treated as if it were 0 .

## Programming Note

Time Base transition events can be used to collect information about activity, as revealed by event counts in PMCs and by addresses in SIAR and SDAR, at periodic intervals.

In multi-threaded systems in which the Time Base registers are synchronized among the threads, Time Base transition events can be used to correlate the Performance Monitor data obtained by the several threads. For this use, software must specify the same TBSEL value for all the threads in the system.
Because the frequency of the Time Base is implementation-dependent, software should invoke a system service program to obtain the frequency before choosing a value for TBSEL.

Time Base Event Enable (TBEE)
0 Time Base transition events are disabled.
1 Time Base transition events are enabled.

## Programming Note

When PMC3 is configured to count the occurrence of Time Base transition events, the events are counted regardless of the value of MMCROTbEE. (See Section 9.4.5.) The occurrence of a Time Base transition causes a Performance Monitor alert only if MMCRO TBEE $=1$.

## BHRB Available (BHRBA)

This field controls whether the BHRB instructions are available in problem state. If an attempt is made to execute a BHRB instruction in problem state when the BHRB instructions are not available, a Facility Unavailable interrupt will occur.
0 clrbhrb and mfbhrbe are not available in problem state.
1 clrbhrb and mfbhrbe are available in problem state unless they have been made unavailable by some other register.
Performance Monitor Event-Based Branch Enable (EBE)
This field controls whether Performance Monitor event-based branches and Performance Monitor event-based exceptions are enabled.
When Performance Monitor event-based branches and exceptions are disabled, no Performance Monitor event-based branches or exceptions occur regardless of the state of BESCR PME .

0 Performance Monitor event-based branches and exceptions are disabled.
1 Performance Monitor event-based branches and exceptions are enabled.

## Programming Note

In order to enable a problem state applications to use the event-based Branch facility for Performance Monitor events, privileged software initializes MMCR1 to specify the events to be counted, and sets MMCR2, and MMCRA to specify additional sampling controls. MMCRO should be initialized with PMCC set to Ob10 or ob11 (to give problem state access to various Performance Monitor registers), PMAE and PMAO set to Os (disabling Performance Monitor alerts), and EBE set to 1 (enabling Performance Monitor event-based branches and exceptions to occur). If the Event-Based Branch facility has not been enabled in the FSCR and HFSCR, it must be enabled in these registers as well.

The above operations by the operating system enable the application to control Performance Monitor event-based branching by means of BESCR PME (to enable or disable Performance Monitor event-based branching) and MMCRO ${ }_{\text {PMAE }}$ (to enable or disable Performance Monitor alerts).

## PMC Control (PMCC)

This field controls whether or not PMCs 5-6 are included in the Performance Monitor, and the accessibility of groups A and B (see Section 9.4.1) of non-privileged SPRs in problem state as described below.

[^16]
## Programming Note

When the PCR makes SPRs unavailable in problem state, they are treated as not implemented, and they are not included in groups $A$ or $B$ regardless of the value of PMCC. Thus when the PCR indicates a version of the architecture prior to V .2 .07 (i.e., $\mathrm{PCR}_{\mathrm{v} 2.06}=1$ ), the PMCC field does not affect SPRs MMCR2 or SIER, which are newly-defined in V. 2.07; these SPRs are treated as unimplemented registers. Accesses to them in problem state result in Hypervisor Emulation Assistance interrupts regardless of the value of PMCC, and Facility Unavailable interrupts do not occur for them. See Section 2.5 for additional information.

00 PMCs 5-6 are included in the Performance Monitor.
Groups A and B are read-only in problem state. If an attempt is made to write to an SPR in group A in problem state, a Hypervisor Emulation Assistance interrupt will occur.
01 PMCs 5-6 are included in the Performance Monitor.
Group A is not allowed to be read or written in problem state, and group $B$ is not allowed to be read in problem state. If an attempt is made, in problem state, to read or write to an SPR in group A, or to read from an SPR in group B, a Facility Unavailable interrupt will occur.
10 PMCs 5-6 are included in the Performance Monitor.
Group A is allowed to be read and written in problem state, and group B except for MMCR1 (SPR 782) is allowed to be read in problem state. If an attempt is made to read MMCR1 in problem state, a Facility Unavailable interrupt will occur.
11 PMCs 5-6 are not included in the Performance Monitor. See Section 9.4.2 for details.
Group A except for PMCs 5-6 (SPRs 775,776 ) is allowed to be read and written in problem state, and group B except for MMCR1 (SPR 782) is allowed to be read in problem state.
If an attempt is made, in problem state, to read or write to PMCs 5-6 (SPRs 775,776 ), or to read from MMCR1, a Facility Unavailable interrupt will occur.
When an SPR is made available by the PMCC field, it is available only if it has not been made unavailable by the HFSCR (see Section 6.2.12).

## Programming Note

In order to give problem state programs the same level of access to the Performance Monitor registers as was specified in Power ISA V 2.06, PMCC must be set to 0b00 (restricting access to read-only) and the PCR should indicate Version 2.06 (restricting access to the set of Performance Monitor SPRs and SPR bits that were defined in V 2.06).

When $\mathrm{PMCC}=0 \mathrm{bOO}$ and a write operation to a Performance Monitor register in group $A$ or $B$ is attempted in problem state, a Hypervisor Emulation Assistance interrupt occurs in order to maintain compatibility with $\vee$ 2.06. For other values of PMCC, write or read operations to group A and read operations from group $B$ that are not allowed result in Facility Unavailable interrupts. Facility Unavailable interrupts provide the operating system with more information about the type of disallowed access that was attempted than the Hypervisor Emulation Assistance interrupt provides. See Section 6.2.11 for additional information.

## Programming Note

In order to prevent applications from accessing Performance Monitor registers, PMCC is set to 0b01.

In order to allow applications limited control over the Performance Monitor, PMCC is set to $0 b 10$ or $0 b 11$. These values are also used when Performance Monitor event-based branches are enabled.

Freeze Counters in Transactional State (FCTS)
0 PMCs are incremented (if permitted by other MMCR bits).
1 PMCs are not incremented when the thread is in Transactional state.

Freeze Counters in Non-Transactional State (FCNTS)
0 PMCs are incremented (if permitted by other MMCR bits).
1 PMCs are not incremented when the thread is in Non-transactional state.
PMC1 Condition Enable (PMC1CE)

This bit controls whether counter negative conditions due to a negative value in PMC1 are enabled.
0 Counter negative conditions for PMC1 are disabled.

1 Counter negative conditions for PMC1 are enabled.

## PMCj Condition Enable (PMCjCE)

This bit controls whether counter negative conditions due to a negative value in any PMCj (i.e., in any PMC except PMC1) are enabled.
0 Counter negative conditions for all PMCjs are disabled.
1 Counter negative conditions for all PMCjs are enabled.
Trigger (TRIGGER)
0 The PMCs are incremented (if permitted by other MMCR bits).
1 PMC1 is incremented (if permitted by other MMCR bits). The PMCjs are not incremented until PMC1 is negative or an enabled condition or event occurs, at which time:
■ the PMCjs resume incrementing (if permitted by other MMCR bits)

- MMCR0 $0_{\text {TRIGGER }}$ is set to 0

See the description of the FCECE bit, above, regarding the interaction between TRIGGER and FCECE.

## - Programming Note

Uses of TRIGGER include the following.
■ Resume counting in the PMCjs when PMC1 becomes negative, without causing a Performance Monitor interrupt. Then freeze all PMCs (and optionally cause a Performance Monitor interrupt) when a PMCj becomes negative. The PMCjs then reflect the events that occurred between the time PMC1 became negative and the time a PMCj becomes negative. This use requires the following MMCRO bit settings.

- TRIGGER=1
- PMC1CE=0
- $\quad$ PMCjCE=1
- TBEE=0
- FCECE=1
- PMAE=1 (if a Performance Monitor interrupt is desired)
- Resume counting in the PMCjs when PMC1 becomes negative, and cause a Performance Monitor interrupt without freezing any PMCs. The PMCjs then reflect the events that occurred between the time PMC1 became negative and the time the interrupt handler reads them. This use requires the following MMCRO bit settings.
- TRIGGER=1
- PMC1CE=1
- TBEE=0
- FCECE=0
- $\quad$ PMAE=1

51 Freeze Counters and BHRB in Problem State Condition (FCPC)

This bit controls the meaning of bit 34 (FCP). See the definition of bit 34 for details.

## Programming Note

In order to enable the FCP bit to freeze counters in problem state regardless of $\mathrm{MSR}_{\mathrm{HV}}, \mathrm{MMCRO}_{\mathrm{FCPC}}$ must be set to 0 .

52 Performance Monitor Alert Qualifier (PMAQ)
This bit provides additional implementation-dependent information about the cause of the Performance Monitor alert. When a Performance Monitor alert occurs, this bit is set to 0 if no additional information is available.53:54

Control Counters 5-6 with Run Latch (CC5-6RUN)
When $\mathrm{MMCRO}_{\text {PMCC }}=\mathrm{b} 11$, the setting of this bit has no effect; otherwise it is defined as follows.

0 PMCs 5 and 6 are incremented if CTRL $_{\text {RUN }}=1$ (if permitted by other MMCR bits).
1 PMCs 5 and 6 are incremented regardless of the value of $C T R L_{\text {RUN }}$ (if permitted by other MMCR bits).

Performance Monitor Alert Occurred (PMAO)
0 A Performance Monitor alert has not occurred since the last time software set this bit to 0 .
1 A Performance Monitor alert has occurred since the last time software set this bit to 0.

This bit is set to 1 by the hardware when a Performance Monitor alert occurs. This bit can be set to 0 only by the $\boldsymbol{m t s p r}$ instruction.

## Programming Note

Software can set this bit to 1 and set PMAE to 0 to simulate the occurrence of a Performance Monitor alert.

Software should set this bit to 0 after handling the Performance Monitor alert.

Freeze Counters in Suspended State (FCSS)
0 PMCs are incremented (if permitted by other MMCR bits).
1 PMCs are not incremented when the thread is in Suspended state.

Freeze Counters 1-4 (FC1-4)
0 PMC1 - PMC4 are incremented (if permitted by other MMCR bits).
1 PMC1 - PMC4 are not incremented.
Freeze Counters 5-6 (FC5-6)
0 PMC5 - PMC6 are incremented (if permitted by other MMCR bits).
1 PMC5-PMC6 are not incremented.
Reserved
Freeze Counters 1-4 in Wait State (FC1-4WAIT)

0 PMCs 1-4 are incremented (if permitted by other MMCR bits).
1 PMCs 1-4, except for PMCs counting events that are not controlled by this bit, are not incremented if $C T R L_{\text {RUN }}=0$.

63 Freeze Counters and BHRB in Hypervisor State (FCH)

0 The PMCs are incremented (if permitted by other MMCR bits) and BHRB entries are written (if permitted by the BHRB Instruction Filtering Mode field in MMCRA).
1 The PMCs are not incremented and BHRB entries are not written if $\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR}=0 \mathrm{~b} 10$.

### 9.4.5 Monitor Mode Control Register 1

Monitor Mode Control Register 1 (MMCR1) is a 64 -bit register as shown below.


## Figure 78. Monitor Mode Control Register 1

MMCR1 enables software to specify the events that are counted by the PMCs.
In the following descriptions, events due to randomly sampled instructions occur only if random sampling is enabled (MMCRA ${ }_{S E}=1$ ); all other events occur whenever the event specification is met regardless of the value of MMCRA SE .

Various events defined below refer to "threshold A" through "threshold H". The table below specifies the number of threshold event counter events corresponding to each of these thresholds.

| Threshold | Events |
| ---: | :--- |
| A | 4096 |
| B | 32 |
| C | 64 |
| D | 128 |
| E | 256 |
| F | 512 |
| G | 1024 |
| H | 2048 |

Table 4: Event Counts for thesholds A-H

The bit definitions of MMCR1 are as follows. Imple-mentation-dependent MMCR1 bits that are not supported are treated as reserved.

## Bit(s) Description

0:31 Problem state access (SPR 782)
Reserved
Privileged access (SPR 782 or 798)
Implementation-dependent

32:39 PMC1 Selector (PMC1SEL)
The value of PMC1SEL specifies the event to be counted by PMC1 as defined below.
All values in the range of E0 - FF that are not specified below are reserved.
Hex Event

00 Disable events. (No events occur.)
01-BF Implementation-dependent
C0-DF Reserved

The following events can occur only when random sampling is enabled (MMCRA $A_{S E}=1$ ). The sampling modes corresponding to each event are listed in parentheses. (The sampling mode is specified in MMCRA ${ }_{S M}$.)
EO The thread has dispatched a randomly sampled instruction. (RIS)
E2 The thread has completed a randomly sampled Branch instruction for which the branch was taken. (RIS, RBS)
E4 The thread has failed to locate a randomly sampled instruction in the primary instruction cache. (RIS)
E6 The threshold event counter has exceeded the number of events corresponding to threshold A (see Table 4). (RIS, RLS, RBS)
E8 The threshold event counter has exceeded the number of events corresponding to threshold E (see Table 4). (RIS, RLS, RBS)
EA The thread filled a block in a data cache with data that were accessed by a randomly sampled Load instruction. (RIS, RLS)
EC The threshold event counter has reached its maximum value. (RIS, RLS, RBS)
The following events can occur regardless of whether random sampling is enabled.

F0 A cycle has occurred. This event is not controlled by MMCRO ${ }_{\text {FC1-4WAIT }}$.
F2 A cycle has occurred in which the thread completed one or more instructions.
F4 The thread has completed a Float-ing-Point, Vector Floating-Point, or VSX Floating-Point instruction other than a

## 1110 Power ISA ${ }^{\text {TM }}$ - Book III

Load or Store instruction to the point at which it has reported all exceptions it will cause.
F6 The thread has failed to locate an ERAT entry during instruction address translation.
F8 A cycle has occurred during which all previously initiated instructions have completed and no instructions are available for initiation.
FA A cycle has occurred during which the RUN bit of the CTRL register for one or more threads of the multi-threaded processor was set to 1.
FC A load type instruction finished. If the instruction caused more than one reference, only one will be counted.
FE The thread has completed an instruction.

## 40:47 PMC2 Selector (PMC2SEL)

The value of PMC2SEL specifies the event to be counted by PMC2 as defined below. All values in the range of EO - FF that are not specified below are reserved.

## Hex Event

00 Disable events. (No events occur.)
01-BF Implementation-dependent
CO-DF Reserved

The following events can occur only when random sampling is enabled (MMCRA ${ }_{S E}=1$ ). The sampling modes corresponding to each event are listed in parentheses. (The sampling mode is specified in MMCRA ${ }_{S M}$.)
E0 The thread has obtained the data for a randomly sampled Load instruction from storage that did not reside in any cache. (RIS, RLS)
E2 The thread has failed to locate the data for a randomly sampled Load instruction in the primary data cache. (RIS, RLS)
E4 The thread filled a block in the primary data cache with data that were accessed by a randomly sampled Load instruction and obtained from a location other than the secondary or tertiary cache. (RIS, RLS)
E6 The threshold event counter has exceeded the number of events corresponding to threshold B (see Table 4). (RIS, RLS, RBS)
E6 The threshold event counter has exceeded the number of events corresponding to threshold $F$ (see Table 4). (RIS, RLS, RBS)
The following events can occur regardless of whether random sampling is enabled.

FO The thread has completed a Store instruction to the point at which it has reported all the exceptions it will cause.
F2 The thread has dispatched an instruction.
F4 A cycle has occurred during which the RUN bit of the thread's CTRL register contained 1.
F6 The thread has failed to locate an ERAT entry during data address translation, and a new ERAT entry corresponding to the data effective address has been written.
F8 An external interrupt for the thread has occurred.
FA The thread has completed a Branch instruction for which the branch was taken.
FC The thread has failed to locate an instruction in the primary cache.
FE The thread has filled a block in the primary data cache with data that were accessed by a Load instruction and obtained from a location other than the secondary cache.

## 48:55

## PMC3Selector (PMC3SEL)

The value of PMC3SEL specifies the event to be counted by PMC3 as defined below.
All values in the range of E0 - FF that are not specified below are reserved.

## Hex Event

00 Disable events. (No events occur.)
01-BF Implementation-dependent
CO-DF Reserved
The following events can occur only when random sampling is enabled (MMCRA ${ }_{S E}=1$ ). The sampling modes corresponding to each event are listed in parentheses. (The sampling mode is specified in MMCRA ${ }_{\text {SM }}$.)
E2 The thread has completed a randomly sampled Store instruction to the point at which it has reported all exceptions it will cause. (RIS,RLS)
E4 The thread has mispredicted either whether or not the branch would be taken, or if taken, the target address of a randomly sampled Branch instruction. (RIS, RBS)
E6 The thread has failed to locate an ERAT entry during data address translation for a randomly sampled instruction. (RIS,RLS)
E8 The threshold event counter has exceeded the number of events corresponding to threshold C (see Table 4). (RIS, RLS, RBS)
EA The threshold event counter has exceeded the number of events corresponding to threshold $G$ (see Table 4). (RIS, RLS, RBS)

The following events can occur regardless of whether random sampling is enabled.

F0 The thread has attempted to store data in the primary data cache but no block corresponding to the real address existed.
F2 The thread has dispatched an instruction.
F4 The thread has completed an instruction when the RUN bit of the CTRL register for all threads on the multi-threaded processor contained 1.
F6 The thread has filled a block in the primary data cache with data that were accessed by a Load instruction.
F8 A Time Base transition event has occurred for the thread. This event is counted regardless of whether or not Time Base transition events are enabled by $\mathrm{MMCRO}_{\text {Tbee }}$.
FA The thread has loaded an instruction from a higher level cache than the tertiary cache.
FC The thread was unable to translate a data virtual address using the TLB.
FE The thread has filled a block in the primary data cache with data that were accessed by a Load instruction and obtained from a location other than the secondary or tertiary cache.

## 56:63 PMC4 Selector (PMC4SEL)

The value of PMC4SEL specifies the event to be counted by PMC4 as defined below.
All values in the range of EO - FF that are not specified below are reserved.
Hex Event
00 Disable events. (No events occur.)
01-BF Implementation-dependent
CO-DF Reserved
The following events can occur only when random sampling is enabled (MMCRA ${ }_{S E}=1$ ). The sampling modes corresponding to each event are listed in parentheses. (The sampling mode is specified in MMCRA ${ }_{S M}$.)
E0 The thread has completed a randomly sampled instruction. (RIS, RLS, RBS)
E4 The thread was unable to translate a data virtual address using the TLB for a randomly sampled instruction. (RIS,RLS)
E6 The thread has loaded a randomly sampled instruction from a higher level cache than the tertiary cache. (RIS)
E8 The thread has filled a block in the primary data cache with data that were accessed by a randomly sampled Load instruction and obtained from a location other than the secondary cache. (RIS, RLS)
EA The threshold event counter has exceeded the number of events corre-
sponding to threshold $D$ (see Table 4). (RIS, RLS, RBS)
EC The threshold event counter has exceeded the number of events corresponding to threshold H (see Table 4). (RIS, RLS, RBS)
The following events can occur regardless of whether random sampling is enabled.

F0 The thread has attempted to load data from the primary data cache but no block corresponding to the real address existed.
F2 A cycle has occurred during which the thread has dispatched one or more instructions.
F4 A cycle has occurred during which the PURR was incremented when the RUN bit of the thread's CTRL register contained 1.
F6 The thread has mispredicted either whether or not the branch would be taken, or if taken, the target address of a Branch instruction.
F8 The thread has discarded prefetched instructions.
FA The thread has completed an instruction when the RUN bit of the thread's CTRL register contained 1.
FC The thread was unable to translate an instruction virtual address using the TLB, and a new TLB entry corresponding to the instruction virtual address has been written.
FE The thread has obtained the data for a Load instruction from storage that did not reside in any cache.
The following event must be supported with an imple-mentation-dependent event code.

The thread has executed an Idmx instruction that accessed a doubleword that contains an effective address within an enabled section of the Load Monitored region. (See Section 3.2.4 of Book I.)
This event only occurs when the Load Monitored Facility and Event-Based Branch Facility are enabled $\left(F_{S C R}^{L M ~ E B B}\right.$ $\left.=0 b 11\right)$.

## Programming Note

This event can be used to measure the frequency of accesses by Idmx to an enabled section of the Load Monitored region. It would typically be used when Load Monitored event-based branch exceptions are disabled (i.e. $B_{E S C R}{ }_{L M E}=0$ ) so that Load Monitored event-based branches do not occur during the measurement process.

## Compatibility Note

In versions of the architecture that precede Version 2.02 the PMC Selector Fields were six bits long, and were split between MMCRO and MMCR1. PMC1-8 were all programmable.

If more programmable PMCs are implemented in the future, additional MMCRs may be defined to cover the additional selectors.

### 9.4.6 Monitor Mode Control Register 2

Monitor Mode Control Register 2 (MMCR2) is a 64-bit register that contains 9 -bit control fields for controlling the operation of PMC1 - PMC6 as shown below.

| C1 | C2 | C3 | C4 | C5 | C6 | Res'd. |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 89 | 1718 | 2627 | 3536 | 4445 | 5354 | 63 |

Figure 79. Monitor Mode Control Register 2
When MMCR0 $_{\text {PMCC }}=0 \mathrm{~b} 11$, fields C1-C4 control the operation of PMC1 - PMC4, respectively and fields C5 and C6 are ignored by the hardware; otherwise, fields C1-C6 control the operation of PMC1-PMC6, respectively. The bit definitions of each Cn field are as follows, where $\mathrm{n}=1, \ldots 6$.

When MMCR0 $0_{\text {PMCC }}$ is set to $0 b 10$ or $0 b 11$, providing problem state programs read/write access to MMCR2, only the FCnP0 bits can be accessed. All other bits are not changed when mtspr is executed in problem state, and all other bits return 0s when mfspr is executed in problem state.

## Bit Description

$0 \quad$ Freeze Counter $n$ in Privileged State (FCnS)

0 PMCn is incremented (if permitted by other MMCR bits).
1 PMCn is not incremented if $\mathrm{MSR}_{\mathrm{HV}} \mathrm{PR}^{2}=0 \mathrm{~b} 00$.
1 Freeze Counter $n$ in Problem State if MSR $_{\boldsymbol{H V}}=\mathbf{O}$ (FCnPO)

0 PMCn is incremented (if permitted by other MMCR bits).
1 PMCn is not incremented if $M_{S R}$ HV PR $=0 b 01$.

## Programming Note

Problem state programs need access to this field in order to enable them to individually enable counters when analyzing sections of code. All the other fields will typically be initialized by the operating system.

2 Freeze Counter $n$ in Problem State if MSR $_{H V}=\mathbf{1}$ (FCnP1)

0 PMCn is incremented (if permitted by other MMCR bits).
1 PMCn is not incremented if $\mathrm{MSR}_{\mathrm{HV} \mathrm{PR}}=0 \mathrm{~b} 11$.

Freeze Counter $n$ in Hypervisor State (FCnH)

0 PMCn is incremented (if permitted by other MMCR bits).
1 PMCn is not incremented if $\mathrm{MSR}_{\mathrm{HV} \mathrm{PR}}=0 \mathrm{~b} 10$.

Bits 54:63 of MMCR2 are reserved.

### 9.4.7 Monitor Mode Control Register A

Monitor Mode Control Register A (MMCRA) is a 64-bit register as shown below.


Figure 80. Monitor Mode Control Register A

MMCRA gives privileged programs the ability to control the sampling process, BHRB filtering, and threshold events.

When MMCRO PMCC $^{\text {is set to } 0 b 10 \text { or Ob11, providing }}$ problem state programs read/write access to MMCRA, the Threshold Event Counter Exponent (TECX) and Threshold Event Counter Multiplier (TECM) fields are read-only, and all other fields return 0s, when mfspr is executed in problem state; all fields are not changed when mtspr is executed in problem state.

## Programming Note

Read/write access is provided to MMCRA in problem state (SPR 770) when MMCR $0_{\text {PMCC }}=0 b 10$ or Ob11 even though no fields can be modified by mtspr because future versions of the architecture may allow various fields of MMCRA to be modified in problem state.

The bit definitions of MMCRA are as follows.

## Bit(s) Description

0:26 Problem state access (SPR 770)
Reserved
Privileged access (SPR 770 or 786)
Implementation-dependent
Bits 27:33 are referred to as the "filtering fields." These fields control the filter criterion used by the hardware when recording Branch instructions in the BHRB. See Section 9.5.1 for specifications of the terminology used and an example of their usage.

## 27 Filter Direct Branch Instructions (FD)

0 Enter direct Branch instructions into the BHRB if allowed by other filtering fields.
1 Do not enter direct Branch instructions into the BHRB.

28 Filter Unconditional Branch Instructions (FU)

0 Enter unconditional Branch instructions into the BHRB if allowed by other filtering fields.
1 Do not enter unconditional Branch instructions into the BHRB.

29 Filter Call Instructions (FC)
0 Enter call instructions into the BHRB if allowed by other filtering fields.
1 Do not enter call instructions into the BHRB.

30

## Filter Return Instructions (FR)

0 Enter return instructions into the BHRB if allowed by other filtering fields.

1 Do not enter return instructions into the BHRB.

BHRB Instruction Filtering Mode (IFM)
00 All taken Branch instructions are entered into the BHRB unless prevented by other
filtering fields.
01 Do not enter jump instructions or return
filtering fields.
01 Do not enter jump instructions or return instructions into the BHRB.

10 Do not enter unconditional Branch instruction addresses into the BHRB, and enter tion addresses into the BHRB, and enter
the target address only if the instruction is indirect.
11 Filter as defined for IFM=10; additionally,
do not enter conditional Branch instruction addresses into the BHRB if $\mathrm{BO}_{0}=1$ or if the "a" bit in the BO field is set to 1 , and enter the target address only if the instruction is indirect.

## Programming Note

Filtering mode 11 provides additional filtering for instructions that provide a hint or for which the outcome does not depend on the value of the CR.

## Filter Jump Instructions (FJ)

0 Enter jump instructions into the BHRB if allowed by other filtering fields.
1 Do not enter jump instructions into the BHRB.

## Threshold Event Counter Exponent

 (TECX)This field species the exponent of the threshold event counter value. See Section 9.4.3 for additional information. The maximum exponent supported is at least 5 .
37 Reserved
Threshold Event Counter Multiplier (TECM)

This field species the multiplier of the threshold event counter value. See Section 9.4.3 for additional information.

## Programming Note

When MMCR0 $_{\text {PMCC }}=0 b 10$ or 0b11, providing problem-state programs read-write access to MMCRA, problem state programs are able to read only the TECX and TECM fields (and are not able to write any fields). The values of these fields are needed during the processing of an event-based branch that occurs due to a counter negative condition for a PMC that was counting "threshold counter exceeded $n$ " events (e.g. MMCR1 $1_{\text {PMC1SEL }}$ $=0 x E 8)$. Reading these fields enables the application to determine the amount by which the threshold was exceeded. Applications are not given access to other fields, and these other fields must initialized by the operating system.

Threshold Event Counter Event (TECE)
This field specifies the event, if any, that is counted by the threshold event counter. The values and meanings are follows.

## Value Event

000 Disable counting.
001 A cycle has occurred.
010 An instruction has completed.
011 Reserved
All other values are implementation-dependent.

## Threshold Start Event (TS)

This field specifies the event that causes the threshold event counter to start counting occurrences of the event specified in the Threshold Event Counter Event (TECE) field. The events only occur if MMCRA $_{\text {SE }}=1$ (random sampling enabled) and one of the sampling modes listed in parenthesis is in effect. (The sampling mode that is currently in effect is specified in MMCRA Sm. $^{\text {.) }}$
0000 Reserved.
0001 The thread has randomly sampled an instruction while it is being decoded. (RIS)
0010 The thread has dispatched a randomly sampled instruction. (RIS)
0011 A randomly sampled instruction has been sent to a facility (e.g. Branch, Fixed Point, etc.) (RIS, RLS, RBS)
0100 The thread has completed a randomly sampled instruction to the point at which it has reported all exceptions it will cause. (RIS, RLS, RBS)
0101 The thread has completed a randomly sampled instruction. (RIS, RLS, RBS)

0110 The thread has failed to locate data for a randomly sampled Load instruction in the primary data cache. (RIS, RLS)
0111 The thread has filled a block in the primary data cache with data that were accessed by a randomly sampled Load instruction. (RIS, RLS)

The definition of the following values depends on whether the access to MMCRA is in problem state or in privileged state.

Problem state access (SPR 770)
1000-1111-Reserved
Privileged access (SPR 770 or 786)
1000-1111-Implementation-dependent

## Threshold End Event (TE)

This field specifies the event that causes the threshold event counter to stop counting occurrences of the event specified in the Threshold Event Counter Event (TECE) field. The events only occur if MMCRA SE $=1$ (random sampling enabled) and one of the sampling modes listed in parenthesis is in effect. (The sampling mode that is currently in effect is specified in MMCRA ${ }_{S M}$.)
0000 Reserved
0001 The thread has randomly sampled an instruction while it is being decoded. (RIS)
0010 The thread has dispatched a randomly sampled instruction. (RIS)
0011 A randomly sampled instruction has been sent to a facility (e.g. Branch, Fixed Point, etc.) (RIS, RLS, RBS)
0100 The thread has completed a randomly sampled instruction to the point at which it has reported all exceptions that it will cause. (RIS, RLS, RBS)
0101 The thread has completed a randomly sampled instruction. (RIS, RLS, RBS)
0110 The thread has failed to locate data for a randomly sampled Load instruction in the primary data cache. (RIS, RLS)
0111 The thread has filled a block in the primary data cache with data that were accessed by a randomly sampled Load instruction. (RIS, RLS)
The definition of the following values depends on whether the access to MMCRA is in problem state or in privileged state.
Problem state access (SPR 770)
1000-1111-Reserved
Privileged access (SPR 770 or 786)
1000-1111-Implementation-dependent
Reserved

## 57:59 Eligibility for Random Sampling (ES)

When random sampling is enabled $\left(\mathrm{MMCRA}_{S E}=1\right)$ and the SM field indicates random instruction sampling (RIS), the encodings of this field specify the instructions that are eligible to be sampled as follows.
000 All instructions
001 All Load and Store instructions
010 All probe no-op instructions
011 Reserved
The definition of the following values depends on whether the access to MMCRA is in problem state or in privileged state.

Problem state access (SPR 770)
100-111-Reserved
Privileged access (SPR 770 or 786)
100-111-Implementation-dependent

When random sampling is enabled (MMCRA $A_{S E}=1$ ) and the SM field indicates random Load/Store Facility sampling (RLS), the encodings of this field specify the instructions that are eligible to be sampled as follows.
000 Instructions for which the thread has attempted to load data from the data cache but no block corresponding to the real address existed.
001 Reserved
010 Reserved
011 Reserved
The definition of the following values depends on whether the access to MMCRA is in problem state or in privileged state.

Problem state access (SPR 770)
100-111-Reserved
Privileged access (SPR 770 or 786)
100-111-Implementation-dependent
When random sampling is enabled (MMCRA ${ }_{S E}=1$ ) and the SM field indicates random Branch Facility sampling (RBS), the encodings of this field specify the instructions that are eligible to be sampled as follows.
000 Instructions for which the thread has either mispredicted whether or not the branch would be taken, or if taken, the target address of a Branch instruction.
001 Instructions for which the thread has mispredicted whether or not the branch of a Branch instruction would be taken because the contents of the Condition Register differed from the predicted contents.
010 Instructions for which the thread has mispredicted the target address of a Branch instruction.

011 All Branch instructions for which the branch was taken.
The definition of the following values depends on whether the access to MMCRA is in problem state or in privileged state.
Problem state access (SPR 770)
100-111-Reserved
Privileged access (SPR 770 or 786)
100-111-Implementation-dependent

60 Reserved
61:62 Random Sampling Mode (SM)
00 Random Instruction Sampling (RIS) Instructions that meet the criterion specified in the ES field for random instruction sampling are eligible to be sampled.
01 Random Load/Store Facility Sampling (RLS) - Instructions that meet the criterion specified in the ES field for random Load/ Store Facility sampling are eligible for sampling.
10 Random Branch Facility Sampling (RBS) - Instructions that meet the criterion specified in the ES field for random Branch Facility sampling are eligible for sampling.
11 Reserved
Random Sampling Enable (SE)
0 Random sampling is disabled.
1 Random sampling is enabled.
See Section 9.4.2.1 for information about random sampling.

### 9.4.8 Sampled Instruction Address Register

The Sampled Instruction Address Register (SIAR) is a 64-bit register.


Figure 81. Sampled Instruction Address Register
When a Performance Monitor alert occurs because of an event caused by execution of a randomly sampled instruction, the SIAR contains the effective address of the instruction if SIER $_{\text {SIARV }}=1$ and contains an undefined value if SIER $_{\text {SIARV }}=0$.
When a Performance Monitor alert occurs because of an event other than an event caused by execution of a randomly sampled instruction, the SIAR contains the effective address of an instruction that was being exe-
cuted, possibly out-of-order, at or around the time that the Performance Monitor alert occurred.

The instruction located at the effective address contained in the SIAR is called the "sampled instruction".

The contents of SIAR may be altered by the hardware if and only if $\mathrm{MMCRO}_{\text {PMAE }}=1$. Thus after the Performance Monitor alert occurs, the contents of SIAR are not altered by the hardware until software sets $\mathrm{MMCRO}_{\text {PMAE }}$ to 1 . After software sets MMCR0 $0_{\text {PMAE }}$ to 1, the contents of SIAR are undefined until the next Performance Monitor alert occurs.

## Programming Note

When the Performance Monitor alert occurs, SIER $_{\text {AMPPR }}$ SAMPHV indicates the value of $M_{\text {HV PR }}$ that was in effect when the sampled instruction was being executed. (The contents of these SIER bits are visible only in privileged state.)

### 9.4.9 Sampled Data Address Register

The Sampled Data Address Register (SDAR) is a 64-bit register.


Figure 82. Sampled Data Address Register
When a Performance Monitor alert occurs because of an event caused by execution of a randomly sampled instruction, the SDAR contains the effective address of the storage operand of the instruction if SIER $_{\text {SDARV }}=1$ and contains an undefined value if SIER $_{\text {SDARV }}=0$.

When a Performance Monitor alert occurs because of an event other than an event caused by execution of a randomly sampled instruction, the SDAR contains the effective address of the storage operand of an instruction that was being executed, possibly out-of-order, at or around the time that the Performance Monitor alert occurred. This storage operand may or may not be the storage operand (if any) of the sampled instruction.

The data located at the effective address contained in the SDAR are called the "sampled data."

The contents of SDAR may be altered by the hardware if and only if $\mathrm{MMCRO}_{\text {PMAE }}=1$. Thus after the Performance Monitor alert occurs, the contents of SDAR are not altered by the hardware until software sets $\mathrm{MMCRO}_{\text {PMAE }}$ to 1 . After software sets MMCR0 $0_{\text {PMAE }}$ to 1, the contents of SDAR are undefined until the next Performance Monitor alert occurs.

### 9.4.10 Sampled Instruction Event Register

The Sampled Instruction Event Register (SIER) is a 64-bit register.

| 0 |
| :--- |
| 0 |

## Figure 83. Sampled Instruction Event Register

When random sampling is enabled and a Performance Monitor alert occurs because of an event caused by execution of a randomly sampled instruction, the SDAR contains information about the sampled instruction. The contents of all fields are valid unless otherwise indicated.

## Programming Note

A Performance Monitor alert occurs because of an event caused by execution of a randomly sampled instruction if random sampling is enabled and a counter negative condition exists in a PMC that was counting events based on randomly sampled instructions.

When random sampling is disabled or when a Performance Monitor alert occurs because of an event that was not caused by execution of a randomly sampled instruction, the contents of the SIER are undefined.
The contents of SIER may be altered by the hardware if and only if $M M C R 0_{\text {PMAE }}=1$. Thus after the Performance Monitor alert occurs, the contents of SIER are not altered by the hardware until software sets MMCR0 $_{\text {PMAE }}$ to 1 . After software sets MMCR0 $0_{\text {PMAE }}$ to 1, the contents of SIER are undefined until the next Performance Monitor alert occurs.
The bit definitions of the SIER are as follows.
0:37 The definition of these bits depends on whether the access to SIER is in problem state or in privileged state.
Problem state access (SPR 768)
Reserved
Privileged access (SPR 768 or 784) Implementation-dependent
38:40 The definition of these bits depends on whether the access to SIER is in problem state or in privileged state.
Problem state access (SPR 768) Reserved

Privileged access (SPR 768 or 784)
38 Sampled MSR PR $^{(S A M P P R)}$
Value of MSR PR when the Performance Monitor alert occurred.

39 Sampled MSR $\boldsymbol{H}_{\boldsymbol{H}}$ (SAMPHV)
Value of $\mathrm{MSR}_{H V}$ when the Performance Monitor alert occurred.
40 Reserved
41 SIAR Valid (SIARV)
Set to 1 when the contents of the SIAR are valid (i.e., they contain the effective address of the sampled instruction); otherwise set to 0 .
SDAR Valid (SDARV)
Set to 1 when the contents of the SDAR are valid (i.e., they contain the effective address of the sampled instruction); otherwise set to 0 .

43 Threshold Exceeded (TE)
Set to 1 by the hardware if the contents of the threshold event counter exceeded the maximum value when the Performance Monitor alert occurred; otherwise set to 0 by the hardware.

## Slew Down

Set to 1 by the hardware if the processor clock was lower than nominal when the Performance Monitor alert occurred; otherwise set to 0 by the hardware.
Slew Up
Set to 1 by the hardware if the processor clock was higher than nominal when the Performance Monitor alert occurred; otherwise set to 0 by the hardware.

46:48 Sampled Instruction Type (SITYPE)
This field indicates the sampled instruction type. The values and their meanings are as follows.

000 The hardware is unable to indicate the sampled instruction type
001 Load Instruction
010 Store instruction
011 Branch Instruction
100 Floating-Point Instruction other than a Load or Store instruction
101 Fixed-Point Instruction other than a Load or Store instruction
110 Condition Register or System Call instruction
111 Reserved

## 49:51 Sampled Instruction Cache Information

 (SICACHE)This field provides cache-related information about the sampled instruction.
000 The hardware is unable to provide any cache-related information for the sampled insttuction.
001 The thread obtained the instruction in the primary instruction cache.

010 The thread obtained the instruction in the secondary cache.
011 The thread obtained the instruction in the tertiary cache.
100 The thread failed to obtain the instruction in the primary, secondary, or tertiary cache
101 Reserved
110 Reserved
111 Reserved
52 Sampled Instruction Taken Branch (SITAKBR)
Set to 1 if the SITYPE field indicates a Branch instruction and the branch was taken; otherwise set to 0 .

53 Sampled Instruction Mispredicted Branch (SIMISPRED)

Set to 1 if the SITYPE field indicates a Branch instruction and the thread has mispredicted either whether or not the branch would be taken, or if taken, the target address; otherwise set to 0 .
54:55 Sampled Branch Instruction Misprediction Information (SIMISPREDI)
If SIMISPRED=1, this field indicates how the thread mispredicted the outcome of a Branch instruction; otherwise this field is set to 0 s.
00 The instruction was not a mispredicted Branch instruction.
01 The thread mispredicted whether or not the branch would be taken because the contents of the Condition Register differed from the predicted contents.
10 The thread mispredicted the target address of the instruction.
11 Reserved
56 Sampled Instruction Data ERAT Miss (SIDERAT)
When the SITYPE field indicates a Load or Store instruction, this field is set to 1 if the thread has failed to locate an ERAT entry during data address translation for the sampled instruction and otherwise is set to 0 .

When the SITYPE field does not indicate a Load or Store instruction, the contents of this field are undefined.

57:59 Sampled Instruction Data Address Translation Information (SIDAXLATE)
This field contains information about data address translation for the sampled instruction. If multiple data address translations were performed, the information pertains to the last translation. The values and their meanings are as follows.

## 1118 Power ISA${ }^{\text {TM }}$ - Book III

000 The instruction did not require data address translation.
001 The thread translated the data virtual address using the TLB.
010 A PTEG required for data address translation for the instruction was obtained from the secondary cache.
011 A PTEG required for data address translation for the instruction was obtained from the tertiary cache.
100 A PTEG required for data address translation for the instruction was obtained from storage that did not reside in any cache.
101 A PTEG required for data address translation for the instruction was obtained from a cache on a different multi-threaded processor that resides on the same chip as the thread.
110 A PTEG required for data address translation for the instruction was obtained from a cache on a different chip from the thread.
111 Reserved

## Sampled Instruction Data Storage Access Information (SIDSAI)

This field contains information about data storage accesses made by the sampled instruction. The values and their meanings are as follows.
000 The instruction did not require data address translation.
001 The instruction was a Read for which the thread obtained the referenced data from the primary data cache.
010 The instruction was a Read for which the thread obtained the referenced data from the secondary cache.
011 The instruction was a Read for which the thread obtained the referenced datafrom the tertiary cache.
100 The instruction was a Read for which the thread obtained the referenced datafrom storage that did not reside in any cache.
101 The instruction was a Read for which the thread obtained the referenced data from a cache on a different multi-threaded processor that resides on the same chip as the thread.
110 The instruction was a Read for which the thread obtained the referenced data from a cache on a different chip from the thread.
111 The instruction was a Store for which the data were placed into a location other than the primary data cache.

## Sampled Instruction Completed (SICMPL)

Set to 1 if the sampled instruction has completed; otherwise set to 0 .

### 9.5 Branch History Rolling Buffer

The Branch History Rolling Buffer (BHRB) is described in Section 2.4 of Book I but only at the level required by application programmers. Additional aspects of the BHRB are described here.

In order to enable problem state programs to use the BHRB, MMCRO $0_{\text {BHRBA }}$ must be set to 1 to enable execution of clrbhrb and mfbhrbe instructions in problem state. Additionally, MMCRO PMCC must be set to 0 b10 or Ob11 to allow problem state programs to read and write the necessary Performance Monitor registers. (See Section 9.4.4.)
If Performance Monitor event-based branching is desired, $\mathrm{MMCRO}_{\text {EBE }}$ must also be set to 1 to enable Performance Monitor event-based branches.

## Programming Note

Enabling Performance Monitor event-based branching eliminates the need for the problem state program to poll MMCRO PMAO in order to determine when a Performance Monitor alert occurs.

The BHRB is written by the hardware if and only if Performance Monitor alerts are enabled by setting $\mathrm{MMCRO}_{\text {PMAE }}$ to 1 . After MMCR0 ${ }_{\text {PMAE }}$ has been set to 1 and a Performance Monitor alert occurs, MMCRO $_{\text {PMAE }}$ is set to 0 and the BHRB is not altered by hardware until software sets $\mathrm{MMCRO}_{\text {PMAE }}$ to 1 again.

When MMCR $_{\text {PMAE }}=1$, mfbhrbe instructions return 0s to the target register.

## Programming Note

mfbhrbe instructions return 0s when MMCR $0_{\text {PMAE }}=1$ in order to prevent software from reading the BHRB while it is being written by hardware.

### 9.5.1 BHRB Filtering

When the BHRB is written by hardware, only those Branch instructions that meet the filtering criterion specified by the filtering fields in $\mathrm{MMCRA}_{27: 33}$ and for which the branch was taken are included.

Filtering restricts the type of Branch instructions that are entered into the BHRB. The filtering criteria are defined using the following terminology.

- Call: A Branch instruction with the LK field set to 1 .
- Return: A bclr instruction with the BH field set to Os.
- Jump: Any Branch instruction that is not a call or a return.
- Conditional Branch: Any Branch instruction other than an I-Form Branch instruction, or a B- or XL-Form Branch instruction with the BO field set to "branch always." (See Figure 41 in Book I.)
- Unconditional Branch: Any Branch instruction other than a conditional branch instruction
■ Indirect Branch: Any XL-Form Branch instruction
- Direct Branch: Any B- or I-Form Branch instruction

Software is able to prevent various combinations of each of the above types of Branch instructions from being entered into the BHRB using the filtering fields in MMCRO and the IFM field in MMCRA. (See Section 9.4.4 and Section 9.4.7.)

## Programming Note

A few examples of how software might use the filtering bits include the following. See Section 9.4.4 for a definition of MMCRO ${ }_{\text {IFM }}$ and Section 9.4 .7 for definitions of filtering bits FD, FU, FC, FR, and FJ. In the following examples, MMCRO ${ }_{\text {IFM }}$ is set to 0 since that filter is not used. Also, the filtering bits listed are set to 1 and other filtering bits are set to Os.

- Enter only return instructions: FC,FJ
- Enter only indirect call instructions: FD, FR,FJ
- Enter only conditional branches: FU
- Enter only indirect jump instructions: FD,FC,FR


### 9.6 Interaction With Other Facilities

If tracing is active $\left(\mathrm{MSR}_{S E}=1\right.$ or $\left.\mathrm{MSR}_{\mathrm{BE}}=1\right)$, the contents of SIAR and SDAR as used by the Performance Monitor facility are undefined and may change even when MMCR0 $0_{\text {PMAE }}=0$.

## Programming Note

A potential combined use of the Trace and Performance Monitor facilities is to trace the control flow of a program and simultaneously count events for that program.

# Chapter 10. Processor Control 

### 10.1 Overview

The Processor Control facility provides a mechanism for the hypervisor to send messages to other threads in
| the system. Privileged non-hypervisor programs are able to send messages to other threads on the same multi-threaded processor; however if the processor is configured into sub-processors, privileged non-hypervisor programs can only send messages to other threads on the same sub-processor.

### 10.2 Programming Model

Both hypervisor-level and privileged-level messages can be sent. Hypervisor-level messages are sent using the msgsnd instruction and cause hypervisor-level
| exceptions when received. Privileged-level messages are sent using the msgsndp instruction and cause privileged-level exceptions when received. For both instructions, the message type and destination threads are specified in a General Purpose Register.

If a message is received by a thread, the exception corresponding to the message type is generated. When the exception is generated, the corresponding interrupt occurs when no higher priority exception exists and the interrupt is enabled ( $\mathrm{MSR}_{\mathrm{EE}}=1$ for the Directed Privileged Doorbell interrupt and $\mathrm{MSR}_{E E}=1$ or $\mathrm{MSR}_{\mathrm{HV}}=0$ for the Directed Hypervisor Doorbell interrupt).

A Directed Privileged Doorbell exception remains until the corresponding interrupt occurs, or the exception is cleared by execution of a mtspr(DPDES) or msgcIrp instruction.

A Directed Hypervisor Doorbell exception remains until the corresponding interrupt occurs, or the exception is
| cleared by execution of a msgcIr instruction.
If a doorbell exception is present and the corresponding interrupt is pended because $\mathrm{MSR}_{\mathrm{EE}}=0$, additional doorbell exceptions are ignored until the exception is cleared.

### 10.3 Processor Control Registers

### 10.3.1 Directed Privileged Doorbell Exception State

The layout of the Directed Privileged Doorbell Exception State (DPDES) register is shown in Figure 84.


Figure 84. Directed Privileged Doorbell Exception State Register

The DPDES register is a 64-bit register. For $\mathrm{t}<\mathrm{T}$, where T is the number of threads on the sub-processor (or on the multi-threaded processor if sub-processors are not supported), bit 63-t corresponds to the thread with privileged thread number t .
The value of bit $t$ indicates the presence of a Directed Hypervisor Doorbell exception on the thread with privileged thread number $t$. Bit $t$ is cleared when a Directed Privileged Doorbell interrupt occurs on thread $t$.
When the contents of DPDES $_{63-\mathrm{t}}$ change from 0 to 1 , a Directed Privileged Doorbell exception will come into existence on privileged thread number $t$ within a reasonable period of time. When the contents of DPDES $_{63}$-t change from 1 to 0 , the existing Directed Privileged Doorbell exception, if any, on privileged thread number $t$, will cease to exist within a reasonable period of time, but not later than the completion of the next context synchronizing instruction or event on privileged thread number t .

The preceding paragraph applies regardless of whether the change in the contents of DPDES $_{63 \text {-t }}$ is the result a msgsndp or msgclrp instruction or of modification of the DPDES register caused by execution of an \| mtspr (DPDES) instruction.

Bits 0:63-T of the DPDES are reserved.

## Version 3.0

I

## Programming Note

The primary use of the DPDES is to provide the means for the hypervisor to save a [sub-]processor's Directed Privileged Doorbell exception state when the set of programs running on the [sub-]processor is swapped out or moved from one [sub-]processor to another. Since there is no such need for a similar function for the hypervisor, there is no similar register for the hypervisor. Privileged programs are able to read the DPDES in order to poll for Directed Privileged Doorbell exceptions when the corresponding interrupt is disabled ( $\mathrm{MSR}_{E E}=1$ ).

### 10.4 Processor Control Instructions

msgsnd, msgsndp, msgcIr, and msgclrp instructions are provided for sending and clearing messages. msgsync is provided to enable the thread that is target of a $\boldsymbol{m s g} \boldsymbol{g} \boldsymbol{n d}$ instruction to ensure that stores performed by the message-sending thread before it executed msg-
snd have been performed with respect to the target thread. msgsndp and msgclrp are privileged instructions, msgsnd, msgclr, and msgsync are hypervisor privileged instructions.


```
msgtype }\leftarrow\operatorname{GPR}(\textrm{RB}\mp@subsup{)}{32:36}{
```

payload $\leftarrow \operatorname{GPR}(\mathrm{RB})_{37: 63}$
If(msgtype $=0 x 05)$ then
send_msg(msgtype, payload)
$\boldsymbol{m} \boldsymbol{s g} \boldsymbol{s} \boldsymbol{n d}$ sends a message to other threads in the system. The message type and destination thread(s) are specified in RB.

## RB



Figure 85. RB Contents for msgsnd
The contents of RB are defined below. Bits 37:63 are referred to as the message payload.

## Field Description

0:31 Reserved
32:36 Type
If Type=0x05, then a Directed Hypervisor Doorbell message is to be sent to the thread(s) specified in the Message Payload field.

All other values of the Type field are reserved; if the instruction is executed with this field set to a reserved value, the instruction is treated as a no-op.

37:38 Broadcast (B)
00 The message is sent to the thread for which PIR $_{44: 63}$ is equal to the value of the PROCIDTAG field in the message payload.
01 The message is sent to all threads on the same sub-processor as the thread for which $\mathrm{PIR}_{44: 63}$ is equal to the value of the PROCIDTAG field in the message payload.

10 The message is sent to all threads on the same multi-threaded processor as the thread for which $\mathrm{PIR}_{44: 63}$ is equal to the value of the PROCIDTAG field in the message payload.
11 Reserved

39:43
44:63

## PROCIDTAG

This field indicates the recipient thread(s) as specified in the B field. If this field set to a value that is not the same as bits PIR $_{44: 63}$ of any thread in the system, then the instruction behaves as if it were a no-op.
The actions taken on receipt of a message are defined in Section 10.2.

This instruction is hypervisor privileged.
Special Registers Altered:
I

## Programming Note

If msgsnd is used to notify the receiver that updates have been made to storage, an [Iw]sync should be placed between the stores and the msgsnd. See Section 5.9.2.

## Version 3.0

## Message Clear

## X-form

msgclr RB

| 31 |  |  | I/I | RB |  | 238 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 |  | 11 | 16 | 21 |
|  |  |  | 31 |  |  |  |

$t \leftarrow$ hypervisor thread number of executing thread
If (msgtype $=0 x 05$ ) then
clear any Directed Hypervisor Doorbell exception
for thread $t$.
$\boldsymbol{m s g} \boldsymbol{l}$ clears a message previously accepted by the thread executing the msgclr.

Let msgtype be (RB) 32: $_{36}$, and let $t$ be the hypervisor thread number of the thread executing the msgclr instruction.

If msgtype $=0 \times 05$, then clear any Directed Hypervisor Doorbell exception that exists on thread t ; otherwise, this instruction is treated as a no-op.

This instruction is hypervisor privileged.
Special Registers Altered:
I
None

- Programming Note
msgclr is typically issued only when $\mathrm{MSR}_{\mathrm{EE}}=0$. If msgelr is executed when $\mathrm{MSR}_{E E}=1$ when a Directed Hypervisor Doorbell interrupt is about to occur, the corresponding interrupt may or may not occur.


## Message Send Privileged X-form

msgsndp RB

| 31 |  | $1 / /$ | RB |  | 142 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

```
msgtype }\leftarrow(\textrm{RB}\mp@subsup{)}{32:36}{
payload }\leftarrow(\textrm{RB}\mp@subsup{)}{37:63}{
t\leftarrow(RB) 57:63
if msgtype = 5 and
    t}\leq\mathrm{ maximum privileged thread number
        on processor or sub-processor
    then
        \mp@subsup{DPDES }{63-t}{*}\leftarrow1
        send_msg(msgtype, payload, t)
```

$\boldsymbol{m s g} \boldsymbol{s} \boldsymbol{n d p}$ sends a message to other threads that are on the same multi-threaded processor (if the processor is not in sub-processor mode) or to other threads that are on the same sub-processor (if the processor is in sub-processor mode). The message type and destination thread(s) are specified in RB.
RB

|  |  | Message Payload |  |  |  |
| :--- | :--- | ---: | ---: | ---: | ---: | ---: |
|  | $/ / /$ | TYPE | $/ / /$ | TIRTAG |  |
| 0 | 32 | 37 | 39 | 57 | 63 |

Figure 86. RB Contents for msgsndp
The contents of RB are defined below. Bits 37:63 are referred to as the message payload.

## Bits Description

37:56 Reserved
57:63 TIRTAG
This message is sent to the thread for which the privileged thread number is equal to contents of the TIRTAG field of the message payload, and one of the following conditions applies.

- for processors that are not partitioned into sub-processors, the thread is sent to the thread on the same multi-threaded processor for which the privileged thread number is equal to the contents of the TIRTAG field of the message payload.
- for processors that are partitioned into sub-processors, the thread is sent to the thread on the same sub-processor for which the privileged thread number is equal to the contents of the TIRTAG field of the message payload.
If $\boldsymbol{m} \boldsymbol{s} \boldsymbol{g} \boldsymbol{s} \boldsymbol{n d} \boldsymbol{p}$ is executed with TIRTAG set to a value greater than the highest privileged thread number on the sub-processor (or on the multi-threaded processor if sub-processors are not supported), then this instruction behaves as a no-op

The actions taken on receipt of a message are defined in Section 10.2.
This instruction is privileged.
Special Registers Altered: DPDES

Programming Note
If $\boldsymbol{m s g} \boldsymbol{s} \boldsymbol{s}$ dp is used to notify the receiver that updates have been made to storage, a Iwsync or sync should be placed between the stores and the msgsndp. See Section 5.9.2.
Message Clear Privileged
msgclrp $\quad$ RB

| 31 | I/I |  | I/I | RB | Rorm |  |  |
| :---: | :---: | :--- | :--- | :--- | :--- | :--- | :--- |
| 0 | 6 |  | 11 | 16 | 21 |  |  |
| 31 |  |  |  |  |  |  |  |

```
msgtype \leftarrow (RB) 32:36
t}\leftarrow\mathrm{ privileged thread number of executing thread
IF(msgtype = 0x05)
    then
        DPDES 
```

$\boldsymbol{m s g} \boldsymbol{c} / r \boldsymbol{p}$ clears a message previously accepted by the thread executing the msgclrp.

Let msgtype be (RB) ${ }_{32: 36}$, and let $t$ be the privileged thread number of the thread executing the msgclrp.
If msgtype $=0 \times 05$, then clear any Directed Privileged Doorbell exception that exists on thread $t$ by setting DPDES $_{63-\mathrm{t}}$ to 0 ; otherwise, this instruction is treated as a no-op.
This instruction is privileged.

## Special Registers Altered:

DPDES

## Programming Note

msgcIrp is typically issued only when $\mathrm{MSR}_{\mathrm{EE}}=0$. If msgcIrp is executed when $\mathrm{MSR}_{\mathrm{EE}}=1$ when a Directed Hypervisor Doorbell interrupt is about to occur, the corresponding interrupt may or may not occur.

Message Synchronize
X-form
msgsync

| 31 | I/I | I/I | I/I |  | 886 | $/$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 |  | 6 | 11 | 16 | 21 |  |
| 31 |  |  |  |  |  |  |

In conjunction with the Synchronize and msgsnd instructions, the msgsync instruction provides an ordering function for stores that have been performed with respect to the thread executing the Synchronize and msgsnd instructions, relative to data accesses by other threads that are performed after a Directed Hypervisor Doorbell interrupt has occurred, as described in the Synchronize instruction description on p. 1021.

This instruction is hypervisor privileged.

## Special Registers Altered:

None

- Programming Note

When used in conjunction with msgsnd, Synchronize with $\mathrm{L}=0$ or 2 is executed on the thread that will execute the msgsnd, and msgsync is executed on another thread -- typically the thread that is the target of the msgsnd, but possibly any other thread (partly because the software that services the Directed Hypervisor Doorbell interrupt may ultimately run on a thread other than that which received the exception). The Synchronize precedes the msgsnd; the msgsync is executed after the Directed Hypervisor Doorbell interrupt occurs, and precedes all instructions that need to "see" the values stored by the stores that are in set A of the memory barrier created by the Synchronize; see Section 5.9.2, "Synchronize Instruction".

# Chapter 11. Synchronization Requirements for Context Alterations 

Changing the contents of certain System Registers, the contents of SLB entries, or the contents of other system resources that control the context in which a program executes can have the side effect of altering the context in which data addresses and instruction addresses are interpreted, and in which instructions are executed and data accesses are performed. For example, changing $\mathrm{MSR}_{\mathrm{IR}}$ from 0 to 1 has the side effect of enabling translation of instruction addresses. These side effects need not occur in program order, and therefore may require explicit synchronization by software. (Program order is defined in Book II.)
An instruction that alters the context in which data addresses or instruction addresses are interpreted, or in which instructions are executed or data accesses are performed, is called a context-altering instruction. This chapter covers all the context-altering instructions. The software synchronization required for them is shown in Table 5 (for data access) and Table 6 (for instruction fetch and execution).
The notation "CSI" in the tables means any context synchronizing instruction (e.g., sc, isync, or rfid). A context synchronizing interrupt (i.e., any interrupt except non-recoverable System Reset or non-recoverable Machine Check) can be used instead of a context synchronizing instruction. If it is, phrases like "the synchronizing instruction", below, should be interpreted as meaning the instruction at which the interrupt occurs. If no software synchronization is required before (after) a context-altering instruction, "the synchronizing instruction before (after) the context-altering instruction" should be interpreted as meaning the context-altering instruction itself.

The synchronizing instruction before the context-altering instruction ensures that all instructions up to and including that synchronizing instruction are fetched and executed in the context that existed before the alteration. The synchronizing instruction after the con-text-altering instruction ensures that all instructions after that synchronizing instruction are fetched and executed in the context established by the alteration. Instructions after the first synchronizing instruction, up to and including the second synchronizing instruction, may be fetched or executed in either context.

If a sequence of instructions contains context-altering instructions and contains no instructions that are affected by any of the context alterations, no software synchronization is required within the sequence.

## Programming Note

Sometimes advantage can be taken of the fact that certain events, such as interrupts, and certain instructions that occur naturally in the program, such as the rfid that returns from an interrupt handler, provide the required synchronization.

No software synchronization is required before or after a context-altering instruction that is also context synchronizing or when altering the MSR in most cases (see the tables). No software synchronization is required before most of the other alterations shown in Table 6, because all instructions preceding the con-text-altering instruction are fetched and decoded before the context-altering instruction is executed (the hardware must determine whether any of these preceding instructions are context synchronizing).
In situations such as context switch in which multiple SPRs are loaded in sequence, it is often the case that the composition of the implicit (implementation-specific, nonarchitectural) synchronizations performed for each individual mtspr will be excessive for the purpose. Software may identify such sequences by placing a $\boldsymbol{m t g s r}$ before the sequence. Hardware may respond to this identification by removing redundant synchronization so that the net synchronization effect approaches that of a single context synchronization at the end of the sequence. A potential side effect of the optimization is that the SPRs specified by the sequence may be loaded in an order other than that specified by the program with the result that if an exception or EBB interrupts the sequence, mtspr instructions past the point of interruption may have loaded their SPRs. When control returns to the interrupted sequence, any such mtspr instructions are re-executed. The programmer must ensure that this side effect will not affect the outcome of the sequence. The degree of optimization is implemen-tation-specific. Transaction failure may compromise optimization.

## Programming Note

Because the individual mtspr instructions in an optimized sequence may be executed in any order, a single sequence should not contain multiple loads of the same SPR. If it does, the final value of the SPR will be indeterminate. Similarly, any set of SPRs for which the relative order of execution of the mtspr instructions in the set matters should not be included in a single optimized sequence.

Unless otherwise stated, the material in this chapter assumes a single-threaded environment.

| Instruction or Event | Required Before | Required After | Notes |
| :---: | :---: | :---: | :---: |
| event-based branch and $\boldsymbol{r f e b b}$ | none | none | 21 |
| interrupt | none | none |  |
| rfid | none | none |  |
| hrfid | none | none |  |
| rfscv | none | none |  |
| sc | none | none |  |
| scv | none | none |  |
| Trap | none | none |  |
| mtmsrd (SF) | none | none | 7 |
| mtmsrd (TS) | none | none |  |
| mtmsrd (TM) | none | none |  |
| $\boldsymbol{m t m s r}[\mathbf{d}]$ (EE) | none | none | 1 |
| $\boldsymbol{m t m s r}[\mathbf{d}]$ (PR) | none | none | 8 |
| $\boldsymbol{m t m s r [ d ] ~ ( F P ) ~}$ | none | none |  |
| $\boldsymbol{m t m s r}[\mathbf{d}](\mathrm{FE} 0, \mathrm{FE} 1)$ | none | none |  |
| $\boldsymbol{m t m s r}[\mathrm{d}]$ (SE, BE) | none | none |  |
| $\boldsymbol{m t m s r [ d ] ~ ( I R ) ~}$ | none | none | 8 |
| $\boldsymbol{m t m s r [ d ] ~ ( R I ) ~}$ | none | none |  |
| mtspr (DEC) | none | none | 9 |
| $\boldsymbol{m t s p r}$ (PIDR) | CSI | CSI | 6 |
| $\boldsymbol{m t s p r}$ (IAMR) | none | CSI |  |
| $\boldsymbol{m t s p r}$ (TFHAR) | none | none |  |
| $\boldsymbol{m t s p r}$ (TEXASR) | none | none |  |
| $\boldsymbol{m t s p r}$ (CTRL) | none | none |  |
| $\boldsymbol{m t s p r}$ (FSCR) | none | CSI |  |
| $\boldsymbol{m t s p r}$ (DPDES) | none | CSI | 17 |
| $\boldsymbol{m t s p r}$ (CIABR) | none | CSI |  |
| $\boldsymbol{m t s p r}$ (HFSCR) | none | CSI |  |
| $\boldsymbol{m t s p r}$ (HDEC) | none | none | 9 |
| $\boldsymbol{m t s p r}$ (HRMOR) | none | CSI | 8,11,17 |
| $\boldsymbol{m t s p r}$ (LPCR) | none | CSI | 11, 12 |
| $\boldsymbol{m t s p r}$ (LPIDR) | CSI | CSI | 6,14,17 |
| $\boldsymbol{m t s p r}$ (PCR) | none | CSI | 17 |
| $\boldsymbol{m t s p r}$ (PTCR) | ptesync | CSI | 3,17 |
| $\boldsymbol{m t s p r}$ (Perf. Mon.) | none | CSI | 15,18 |
| $\boldsymbol{m t s p r}$ (BESCR) | none | CSI | 16,18 |
| slbie | none | CSI |  |
| slbieg | none | CSI | 4,6 |
| slbia | none | CSI | 4 |
| slbmte | none | CSI | 4,8,10 |
| tlbie | none | CSI | 4,6 |
| tlbiel | none | CSI | 4 |
| Store(PTE) | none | \{ptesync, CSI\} | 5,6,8 |
| Store(STE) | none | \{ptesync, CSI\} | 5,6,8 |
| Store(PRTE) | none | \{ptesync, CSI $\}$ | 5,6,8 |

Table 6: Synchronization requirements for instruction fetch and/or execution

| Instruction or <br> Event | Required <br> Before | Required <br> After | Notes |
| :--- | :--- | :--- | :--- |
| Store(PATE) | none | \{ptesync, <br> CSI $\}$ | $5,6,8$ |
| none | 21 |  |  |
| transaction failure <br> and all TM <br> instructions <br> except tcheck | none | no |  |

Table 6: Synchronization requirements for instruction fetch and/or execution

## Notes:

1. The effect of changing the EE bit is immediate, even if the $\boldsymbol{m t m s r}[d]$ instruction is not context synchronizing (i.e., even if $\mathrm{L}=1$ ).

- If an mtmsr[d] instruction sets the EE bit to 0 , neither an External interrupt, a Decrementer interrupt nor a Performance Monitor interrupt occurs after the mtmsr$[d]$ is executed.
- If an mtmsr[d] instruction changes the EE bit from 0 to 1 when an External, Decrementer, Performance Monitor or higher priority exception exists, the corresponding interrupt occurs immediately after the $\boldsymbol{m t m s r}[\boldsymbol{d}]$ is executed, and before the next instruction is executed in the program that set EE to 1.
- If a hypervisor executes the mtmsr[d] instruction that sets the EE bit to 0, a Hypervisor Decrementer interrupt does not occur after mtmsr[d] is executed as long as the thread remains in hypervisor state.
- If the hypervisor executes an mtmsr[d] instruction that changes the EE bit from 0 to 1 when a Hypervisor Decrementer or higher priority exception exists, the corresponding interrupt occurs immediately after the mtmsr[d] instruction is executed, and before the next instruction is executed, provided HDICE is 1.

2. Synchronization requirements for this instruction are implementation-dependent.
3. The PTCR controls all implicit and explicit storage accesses performed by all threads on the processor when translation is enabled. Modifying the PTCR requires that the following conditions be achieved on all threads on the processor.

- translation is disabled
- all previous accesses (implicit and explicit) initiated with translation enabled have been performed with respect to all threads
■ no subsequent accesses which require translation have been initiated

4. For data accesses, the context synchronizing instruction before the slbie, slbieg, slbia, slbmte, tlbie, or tlbiel instruction ensures that all preceding instructions that access data storage have completed to a point at which they have reported all exceptions they will cause.
The context synchronizing instruction after the slbie, slbieg, slbia, slbmte, tlbie or tlbiel instruction ensures that storage accesses associated with instructions following the context synchronizing instruction will not use the TLB entry(s) being invalidated.
I
(For tlbie and tlbiel, if it is necessary to order stor- age accesses associated with preceding instructions, or Reference and Change bit updates associated with preceding address translations, with respect to subsequent data accesses, a pte-
sync instruction must also be used, either before or after the tlbie or tlbiel instruction. These effects of the ptesync instruction are described in the last paragraph of Note 5.)
5. The notation "\{ptesync,CSI\}" denotes an instruction sequence. Other instructions may be interleaved with this sequence, but these instructions must appear in the order shown.

No software synchronization is required before the Store instruction because (a) stores are not performed out-of-order and (b) address translations associated with instructions preceding the Store instruction are not performed again after the store has been performed (see Section 5.5). These properties ensure that all address translations associated with instructions preceding the Store instruction will be performed using the old contents of the PTE.

The ptesync instruction after the Store instruction ensures that all searches of the Page Table that are performed after the ptesync instruction completes will use the value stored (or a value stored subsequently). The context synchronizing instruction after the ptesync instruction ensures that any address translations associated with instructions following the context synchronizing instruction that were performed using the old contents of the PTE will be discarded, with the result that these address translations will be performed again and, if there is no corresponding entry in any TLB, SLB, page walk cache, cache of Partition or Process Table entries, or implementation-specific address translation lookaside information, will use the value stored (or a value stored subsequently).
The ptesync instruction also ensures that all storage accesses associated with instructions preceding the ptesync instruction, and all Reference and Change bit updates associated with additional address translations that were performed, by the thread executing the ptesync instruction, before the ptesync instruction is executed, will be performed with respect to any thread or mechanism, to the extent required by the associated Memory Coherence Required attributes, before any data accesses caused by instructions following the ptesync instruction are performed with respect to that thread or mechanism.
6. There are additional software synchronization requirements for this instruction in multi-threaded environments (e.g., it may be necessary to invalidate one or more TLB entries on all threads in the system and to be able to determine that the invalidations have completed and that all side effects of the invalidations have taken effect).

Section 5.10 gives examples of using tlbie, Store, and related instructions to maintain the Page Table, in both multi-threaded environments and
environments consisting of only a single-threaded processor.

## Programming Note

In a multi-threaded system, if software locking is used to help ensure that the requirements described in Section 5.10 are satisfied, the Iwsync instruction near the end of the lock acquisition sequence (see Section B.2.1.1 of Book II) may naturally provide the context synchronization that is required before the alteration.
7. The alteration must not cause an implicit branch in effective address space. Thus, when changing $\mathrm{MSR}_{\mathrm{SF}}$ from 1 to 0 , the mtmsrd instruction must have an effective address that is less than $2^{32}-4$. Furthermore, when changing $M_{S R}$ from 0 to 1 , the mtmsrd instruction must not be at effective address $2^{32}-4$ (see Section 5.3.2 on page 981).
8. The alteration must not cause an implicit branch in real address space. Thus the real address of the context-altering instruction and of each subsequent instruction, up to and including the next context synchronizing instruction, must be independent of whether the alteration has taken effect.

## Programming Note

If it is desired to set $\mathrm{MSR}_{\mathrm{IR}}$ to 1 early in an operating system interrupt handler, advantage can sometimes be taken of the fact that $\mathrm{EA}_{0: 3}$ are ignored when forming the real address when address translation is disabled and $\mathrm{MSR}_{\mathrm{HV}}=0$. For example, if address translation resources are set such that effective address 0xn000_0000_0000_0000 maps to real address 0x000_0000_0000_0000 when address translation is enabled, where n is an arbitrary 4-bit value, the following code sequence, in real page 0 , can be used early in the interrupt handler.

```
    la rx,target
    li ry,0xn000
    sldi ry,ry,48
    or rx,rx,ry # set high-order
                                    nibble of target
                                    addr to 0xn
    mtctr rx
    bcctr # branch to targ
targ: mfmsr rx
    orir x,rx,0x0020
    mtmsrd rx # set MSR[IR] to 1
```

The mtmsrd does not cause an implicit branch in real address space because the real address of the next sequential instruction is independent of $\mathrm{MSR}_{\mathrm{IR}}$. Using mtmsrd, rather than rfid (or similar context synchronizing instruction that alters the control flow), may yield better performance on some implementations.
(Variations on the technique are possible. For example, the target instruction of the bcctr can be in arbitrary real page $P$, where $P$ is a 48 -bit value, provided that effective address 0xn || P || 0x000 maps to real address P \| 0x000 when address translation is enabled.)
9. The elapsed time between the contents of the Decrementer or Hypervisor Decrementer becoming negative and the signaling of the corresponding exception is not defined.
10. If an slbmte instruction alters the mapping, or associated attributes, of a currently mapped ESID, the slbmte must be preceded by an slbie (or slbia) instruction that invalidates the existing translation. This applies even if the corresponding entry is no longer in the SLB (the translation may still be in implementation-specific address translation lookaside information). No software synchronization is needed between the slbie and the slbmte, regardless of whether the index of the SLB entry (if any) containing the current translation is the same as the SLB index specified by the slbmte.
No slbie (or slbia) is needed if the slbmte instruction replaces a valid SLB entry with a mapping of a
different ESID (e.g., to satisfy an SLB miss). However, the slbie is needed later if and when the translation that was contained in the replaced SLB entry is to be invalidated.
11. When the HRMOR or the VC field of the LPCR is modified, software must invalidate all implementa-tion-specific lookaside information used in address translation that depends on the old contents of the register or field (i.e., the contents immediately before the modification). The slbia instruction can be used to invalidate all such implementation-specific lookaside information.
12. A context synchronizing instruction or event that is executed or occurs when LPCR $_{\text {MER }}=1$ does not necessarily ensure that the exception effects of $\mathrm{LPCR}_{\text {MER }}$ are consistent with the contents of LPCR MER. See Section 2.2.
13. This line applies regardless of which SPR number (13 or 29) is used for the AMR.
14. LPIDR must not be altered when $M S R_{D R}=1$ or $M S R_{I R}=1$; if it is, the results are undefined.

## Programming Note

Altering LPIDR when $\mathrm{MSR}_{\mathrm{IR}=1 \text { or } \mathrm{MSR}_{\mathrm{DR}}=1 \text { is }}$ prohibited because of the difficulty of avoiding an implicit branch relative to the value of enabling software to avoid using hypervisor real addressing mode for the operation. (The tables used for translation are determined by the partition ID. See Section 5.7.6 for details.)
15. This line applies to the following Performance Monitor SPRs: PMC1-6, MMCR0, MMCR1, MMCR2, and MMCRA.
16. This line applies to all SPR numbers that access the BESCR $(800-803,806)$.
17. There are additional software synchronization requirements when an mtspr instruction modifies this SPR in a multi-threaded environment. See Section 2.7.
18. As an alternative to a CSI, the execution of an rfebb instruction or the occurrence of an event-based branch is sufficient to provide the necessary synchronization.
19. These instructions and events, with the exception of nested tbegin. nested tend., TM instructions that except or are described to be treated as noops, Transaction Abort Conditional instructions that do not abort, and events and rfebb instructions for which the event did not take place in Transactional state, will change MSR Tss . No software synchronization is required.

## Power ISA Book I-III Appendices

## Appendix A. Illegal Instructions

With the exception of the instruction consisting entirely of binary $0 s$, the instructions in this class are available for future extensions of the Power ISA; that is, some future version of the Power ISA may define any of these instructions to perform new functions.
The following primary opcodes are illegal.
1, 5, 6
The following primary opcodes have unused extended opcodes. Their unused extended opcodes can be determined from the opcode maps in Appendix $C$ of Book Appendices. All unused extended opcodes are illegal.
$4,19,30,31,56,5,58,59,60,62,63$
An instruction consisting entirely of binary 0 s is illegal, and is guaranteed to be illegal in all future versions of this architecture.

The following instruction is illegal in privileged state:

- Idmx


## Appendix B. Reserved Instructions

The instructions in this class are allocated to specific purposes that are outside the scope of the Power ISA.

The following types of instruction are included in this class.

1. The instruction having primary opcode 0 , except the instruction consisting entirely of binary 0 s (which is an illegal instruction; see Section 1.8.2, "Illegal Instruction Class" on page 22) and the extended opcode shown below.

256 Service Processor "Attention"
2. Instructions for the POWER Architecture that have not been included in the Power ISA. These are listed in Section A.31, "Discontinued Opcodes" and Section A.33.1, "Discontinued Opcodes".
3. Implementation-specific instructions used to conform to the Power ISA specification.
4. Any other implementation-dependent instructions that are not defined in the Power ISA.

## Appendix C. Opcode Maps

This appendix contains opcode maps showing the primary opcodes, extended opcodes, and expanded opcodes.

Table 7 describes the conventions used in the opcode maps.
The instruction consisting entirely of binary 0s causes the system illegal instruction error handler to be
invoked for all members of the POWER family, and this is likely to remain true in future models (it is guaranteed in the Power ISA). An instruction having primary opcode 0 but not consisting entirely of binary 0 s is reserved except for the following extended opcode (instruction bits 21:30).

256 Service Processor "Attention"

Table 7: Opcode Maps Legend

| $\begin{array}{\|l\|l} \hline \text { po } \\ \text { version } \end{array}$ | mnemonicprivilegebook <br> format | po primary opcode (decimal format) |
| :---: | :---: | :---: |
| xop version | mnemonic  <br> privilege book <br> format | xop extended or expanded opcode image (binary format) |
|  |  | 0 instruction bit corresponding to an extended/expanded opcode bit having value of 0 <br> 1 instruction bit corresponding to an extended/expanded opcode bit having value of 1 <br> / reserved instruction bit, must have value of 0 , otherwise invalid form instruction bit corresponding to an operand or control bit, can have a value of either 0 or 1 |
|  |  | book |
|  |  | Book instruction defined |
|  |  | version |
|  |  | ISA version instruction introduced |
|  |  | privilege |
|  |  | P privileged instruction |
|  |  | H hypervisor-privileged instruction |
|  |  | format |
|  |  | instruction format |



Illegal opcode
Opcode having no previous or current assignment, available for future use
Defined opcode (primary, extended, or expanded)
Opcode assigned to a defined instruction
Primary opcode having an extended opcode field Opcode having extended opcode field used to identify multiple instructions
Extended opcode having an expanded opcode field Opcode having expanded opcode field used to identify multiple instructions
Reserved opcode (primary, extended, or expanded) Opcode is not available for future use without careful consideration

1. Opcode corresponds to an instruction defined in a previous version of the architecture that has been subsequently removed from the architecture. The opcode is treated as an illegal opcode.
2. Or, opcode is reserved for implementation-dependent use.

These opcodes will not be assigned a meaning in the Power ISA except after careful consideration of the effect of such assignment on existing implementations.
$\square$ Invalid form opcode
Opcode corresponding to a defined instruction encoding with one or more reserved opcode bits having a value of 1

Table 8: Primary Opcode Map (opcode bits 0:5)


Table 9: EXT17: Extended Opcode Map for Primary Opcode 17 (opcode bits 27:30)

| 00 | 01 |  | 10 |  | 11 |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | $\begin{array}{ll} 01 & \text { scv } \\ \hline \text { v3.0 } \end{array}$ | $\begin{array}{r} \mathrm{I} \\ \left.\mathrm{sc}\right\|_{\mathrm{PPC}} ^{1 / 2} \end{array}$ | sc | $\begin{gathered} \mathrm{I}[1 / \\ \mathrm{sc}\{\text { \{invalid\} } \end{gathered}$ | sc |
| 00 | 01 |  | 10 |  | 11 |

Table 10:EXT30: Extended Opcode Map for Primary Opcode 30 (opcode bits 27:30)


Table 11:EXT57: Extended Opcode Map for Primary Opcode 57 (opcode bits 30:31)

| 00 |  | 01 | 10 | 11 |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll} 00 & \text { Ifdp } \\ \mathrm{v} 2.05 & \\ \hline \end{array}$ | I | \{reserved\} | $\begin{array}{\|ll} \hline 10 & \text { lxsd } \\ \hline \text { v3.0 } \end{array}$ | $\begin{array}{r\|l\|l} \hline \text { I } & 11 \\ \text { DS } & \text { v3.0 } \\ \hline \end{array}$ | Ixssp | DS |
| 00 |  | 01 | 10 |  | 11 |  |

Table 12:EXT58: Extended Opcode Map for Primary Opcode 58 (opcode bits 30:31)


Table 13:EXT61: Extended Opcode Map for Primary Opcode 61 (opcode bits 29:31)


Table 14:EXT62: Extended Opcode Map for Primary Opcode 62 (opcode bits 30:31)

| 00 |  | 01 |  | 10 |  | 11 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{\|ll} \hline 00 & \text { std } \\ \hline \text { PPC } & \\ \hline \end{array}$ |  | stdu | $\begin{gathered} \hline \mathrm{I} \mid 10 \\ \mathrm{DS} \end{gathered} \mathrm{I}_{\mathrm{V} .03}$ | stq |  |  |
| 00 |  | 01 |  | 10 |  | 11 |

Table 15:EXT04: Extended Opcode Map for Primary Opcode 4 (opcode bits 21:31) (Sheet 1 of 8)


Version 3.0

Table 15:EXT04: Extended Opcode Map for Primary Opcode 4 (opcode bits 21:31) (Sheet 2 of 8 )

|  | 001000 | 001001 | 001010 | 001011 | 001100 | 001101 | 001110 | 001111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 |  |  | $\left\|\begin{array}{c} \text { 00000 001010 } \\ \text { v2.03 } \\ \text { vaddfp } \end{array} \quad \begin{array}{r} \text { Ix } \end{array}\right\|$ |  | $\left\|\begin{array}{lr}00000 \\ \text { v2.00 } 011100 \\ \text { vmrghb } & \mathrm{I} \\ & v x\end{array}\right\|$ |  | $\left\|\begin{array}{lr}00000001110 \\ \text { v2.03 } \\ \text { vpkuhum } & \mathrm{I} \\ & \mathrm{vx}\end{array}\right\|$ |  | 00000 |
| 00001 |  |  | $\left\lvert\, \begin{array}{lr} 00001 & 001010 \\ \text { v2.03 } & \text { vsubfp } \\ \text { vx } & \text { I } \end{array}\right.$ |  | $\left\|\begin{array}{rr} 0.0001 & \text { 001100 } \\ \text { v2.03 } & \text { vmrghh } \\ \text { vx } \end{array}\right\|$ |  |  |  | 00001 |
| 00010 | $\|$20010 001000 I <br> v2.07 vmulouw $v x$ | 00010 001001  <br> v2.07  <br>   <br>  I I. <br>  vx |  |  |  |  | $\left\|\begin{array}{lr}0.0010 & 001110 \\ \text { v2.03 } & \text { vkuhus } \\ & \mathrm{vx}\end{array}\right\|$ |  | 00010 |
| 00011 |  |  |  |  |  |  | 00011  <br> v201110  <br> vpkuwus I <br>  vx$\|$ |  | 00011 |
| 00100 |  |  | $\left\|\begin{array}{lrr}00100 & 001010 \\ \text { v2.03 } & \text { vrefp } & \mathrm{I} \\ & & v x\end{array}\right\|$ |  | $\begin{array}{\|cr\|} \hline 00100 \\ \mathbf{v}_{2}^{0001100} & \text { vmrglb } \\ \hline \end{array}$ |  | 00100 001110 <br> v2.03  vpkshus <br>  IX$\|$ |  | 00100 |
| 00101 | $\begin{array}{\|lr} \hline 00101 \text { ve1000 } & \text { I } \\ \text { v2.03 } & \\ & v x \\ \hline \end{array}$ |  | $\|$$001010{ }^{001010}$ vis <br> v2.03 vrsqrtefp <br>  $v x$ |  | $\begin{array}{\|lr\|} \hline 00101 & \text { 001100 } \\ \text { v2.03 } & \text { vmrglh } \\ & \mathrm{vx} \\ \hline \end{array}$ |  | 00101 001110 <br> v2.03 vpkswus <br>  vx$\|$ |  | 00101 |
| 00110 | $\left\lvert\, \begin{array}{\|cc\|} \hline \text { O0110 ve1000 } & \text { I } \\ \text { v2.07 } & \text { vmulosw } \\ & v x \end{array}\right.$ |  | $\left\|\begin{array}{rr} 00110 & \text { 001010 } \\ \text { v2.03 } & \text { vexptefp } \\ & \mathrm{vx} \end{array}\right\|$ |  | $\left\lvert\, \begin{array}{\|lr} \hline 00110 & \text { 001100 } \\ \text { v2.03 } & \text { vmrglw } \\ & \mathrm{Ix} \\ \hline \end{array}\right.$ |  | $\left\|\begin{array}{lr}00110 & 001110 \\ \text { v2.03 }\end{array} \quad \begin{array}{r}\text { vpkshss } \\ \\ v x\end{array}\right\|$ |  | 00110 |
| 00111 |  |  | $\|$00111 001010 $^{\text {vole }}$ I <br> v2.03  |  |  |  | $\left\|\begin{array}{lr}00111 & \text { 001110 } \\ \text { v2.03 }\end{array} \quad \begin{array}{r}\text { I } \\ \\ \end{array}\right\|$ |  | 00111 |
| 01000 | $\|$01000 001000 I <br> v2.03  <br>  $v x$ |  | $\left\|\begin{array}{cc} 01000 & 001010 \\ \text { v2.03 } \end{array} \quad \begin{array}{r} \text { vrfin } \\ \end{array}\right\|$ |  | 01000 001100 <br> v2.03 vspltb <br> In vx | $\begin{aligned} & \mathrm{T} \\ & \times \begin{array}{l} 01000 \\ \text { vextrat } \\ \text { vextractub } \end{array} \\ & \hline \end{aligned}$ |  |  | 01000 |
| 01001 | $\left\lvert\, \begin{array}{\|cr\|} \hline 01001 \text { ve11000 } & \text { I } \\ \text { v2.03 } & \\ \hline \end{array}\right.$ |  | $\begin{array}{\|c\|c\|} \hline 01001 & 001010 \\ \text { v2.03 } & \text { vrfiz } \\ & \text { vx } \\ \hline \end{array}$ |  |  |  |  |  | 01001 |
| 01010 | $\left\lvert\, \begin{array}{\|cc\|} \hline 01010 \text { ve1000 } & \text { I } \\ \text { v2.07 } & \text { vmuleuw } \\ & v x \end{array}\right.$ |  | $\left\lvert\, \begin{array}{ll} 01010 & 001010 \\ \text { v2.03 } & \text { vrfip } \\ & \text { vx } \end{array}\right.$ |  |  |  | $\left.\begin{array}{r\|rr} \hline \text { I } \\ \text { vx } \left\lvert\, \begin{array}{rl} 01010 & 001110 \\ \text { v2.03 } \end{array}\right. & \text { vupklsb } & \mathrm{I} \\ & \mathrm{vx} \end{array} \right\rvert\,$ |  | 01010 |
| 01011 |  |  | $\begin{array}{lr} 01011 & \text { 001010 } \\ \text { v2.03 } & \text { vrfim } \\ & \mathrm{vx} \end{array}$ |  |  | $\left.\right\|^{01011} \begin{aligned} & 001101 \\ & \text { v3.0 } \end{aligned} \text { vextractd }$ |  |  | 01011 |
| 01100 | $\left\lvert\, \begin{array}{\|lr} \hline 01100 \text { ve1000 } & \text { I } \\ \text { v2.03 } \end{array}\right.$ |  |  |  |  |  |  |  | 01100 |
| 01101 | $\begin{array}{\|lr\|} \hline 01101 \text { ve1000 } & \text { I } \\ \hline \text { v2.03 } & \\ \hline \text { vmulesh } & v x \\ \hline \end{array}$ |  | $\begin{array}{\|l\|} \hline 01101 \\ \hline 001010 \\ \text { v2.03 } \end{array} \quad \text { vcfsx } \left.\begin{aligned} & \text { I } \\ & v x \end{aligned} \right\rvert\,$ |  |  |  |  |  | 01101 |
| 01110 | $\left\|\begin{array}{\|cc} 01110 & \text { ve1000 } \\ \text { v2.07 } & \text { vmulesw } \\ & \\ \text { vx } \end{array}\right\|$ |  | $\left\|\begin{array}{\|cr} 01110 \\ 0001010 & \text { I } \\ \text { v2.03 } & \text { vctuxs } \\ & v x \end{array}\right\|$ |  |  |  | $\begin{gathered} \mathrm{I} \\ \mathrm{vx} \end{gathered}$ |  | 01110 |
| 01111 |  |  | $\left\|\begin{array}{rr} 01111 & 001010 \\ \text { v2.03 } \end{array} \begin{array}{rrr} \text { vctsxs } & \mathrm{I} \\ & \mathrm{vx} \end{array}\right\|$ |  |  | $\left\lvert\, \begin{gathered} 01111 \\ \begin{array}{c} 0111 \\ \text { v3.0 } \end{array} \\ \text { vinsertd } \end{gathered}\right.$ | $\left.\begin{array}{rrr} \mathrm{I} \\ \mathrm{vx} \left\lvert\, \begin{array}{c} 01111 \\ \text { v2.03 } \end{array}\right. & \begin{array}{c} \text { vupkl110 } \\ \text { vupx } \end{array} \\ \mathrm{vx} \end{array} \right\rvert\,$ |  | 01111 |
| 10000 | $\left\|\begin{array}{lr}10000 & 001000 \\ \text { v2.07 } & \text { vpmsumb } \\ & \mathrm{IX}\end{array}\right\|$ |  | $\left\|\begin{array}{cc} 10000 \\ \text { vol1010 } \\ \text { vmaxfp } \\ & \mathrm{vx} \end{array}\right\|$ |  | 10000 001100  <br> v2.03 vslo I <br>   vx |  |  |  | 10000 |
| 10001 | $\left\|\begin{array}{\|cr\|}\hline 10001 \text { v01000 } & \text { I } \\ \text { v2.07 } & \text { vpmsumh } \\ & v x\end{array}\right\|$ |  | $\left\|\begin{array}{lr}10001 \\ \text { v201010 } \\ \text { vminfp } & \text { I } \\ & v x\end{array}\right\|$ |  | 10001 001100 I <br> v2.03 vsro vx <br>    |  | 10001 001110  <br> vpkudum I <br>  vx$\|$ |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  | 10011001110  <br> v2.07 vpkudus <br>  vx <br> v.  |  | 10011 |
| 10100 | $\left\lvert\, \begin{array}{\|cc\|} \hline 10100 & 001000 \\ \text { v2.07 } & \text { vcipher } \\ \text { va } \end{array}\right.$ |  |  |  | $\left\|\begin{array}{\|cc\|} \hline 10100 & 001100 \\ & \text { I } \\ \text { v2.07 } & \text { vgbbd } \end{array} \quad \mathrm{vx}\right\|$ |  |  |  | 10100 |
| 10101 | $\|$10101001000  <br> vncipher I <br> v2.07   <br>  $v x$ |  |  |  |  |  | 10101 001110 <br> v2.07 vpksdus <br>  vx$\|$ |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  | 10110 |
| 10111 | $\|$  <br> 10111 0010000 <br> v2.07 vsbox <br>  vx |  |  |  | $\left\|\begin{array}{lr}\hline 10111 & \text { 0011100 } \\ \text { v3.0 } & \text { vbpermd } \\ & \mathrm{vx}\end{array}\right\|$ |  | $\left\|\begin{array}{lr}10111 \text { 001110 } & \mathrm{I} \\ \text { v2.07 } & \text { vksdss } \\ & \mathrm{vx}\end{array}\right\|$ |  | 10111 |
| 11000 |  |  |  |  |  | $\begin{array}{\|l\|l} 11000 & 001101 \\ & \text { vextublx } \\ \text { v3.0 } \end{array}$ |  |  | 11000 |
| 11001 |  |  |  |  |  | $\left\lvert\, \begin{array}{cc} 11001 & 001101 \\ \text { v3.0 } & \text { vextuhlx } \\ \hline \end{array}\right.$ |  |  | 11001 |
| 11010 | $\qquad$ |  |  |  | $\begin{array}{\|cc\|} \hline 11010 & \text { 001100 } \\ \text { v2.07 } & \text { vmrgow } \\ & \mathrm{I} \\ \hline \end{array}$ |  | $\begin{array}{r} \mathrm{I} \\ \mathrm{vx} \end{array}$ |  | 11010 |
| 11011 |  |  |  |  |  |  | $\left\|\begin{array}{lr}11011 \\ \text { v201110 } \\ \text { vupklsw } & \text { I } \\ & v x\end{array}\right\|$ |  | 11011 |
| 11100 |  |  |  |  |  | $\begin{array}{\|l\|l\|} \hline 11100 & 001101 \\ \text { vex.0 } \\ \text { vextubrx } \\ \hline \end{array}$ |  |  | 11100 |
| 11101 |  |  |  |  |  | $\begin{array}{\|l\|l\|} \hline 11101 & 001101 \\ \text { v3.0 } & \text { vextuhrx } \\ \hline \end{array}$ |  |  | 11101 |
| 11110 | $\|$  <br> 11110 001000 <br> vsumsws I <br> v2.03   <br>   <br>   |  |  |  | 11110 001100 <br> vmrgew  <br> v2.07  I <br>   |  |  |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  | 11111 |
|  | 001000 | 001001 | 001010 | 001011 | 001100 | 001101 | 001110 | 001111 |  |

Table 15:EXT04: Extended Opcode Map for Primary Opcode 4 (opcode bits 21:31) (Sheet 3 of 8)

|  | 010000 | 010001 | 010010 | 010011 | 010100 | 010101 | 010110 | 010111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 |  |  |  |  |  |  |  |  | 00000 |
| 00001 |  |  |  |  |  |  |  |  | 00001 |
| 00010 |  |  |  |  |  |  |  |  | 00010 |
| 00011 |  |  |  |  |  |  |  |  | 00011 |
| 00100 |  |  |  |  |  |  |  |  | 00100 |
| 00101 |  |  |  |  |  |  |  |  | 00101 |
| 00110 |  |  |  |  |  |  |  |  | 00110 |
| 00111 |  |  |  |  |  |  |  |  | 00111 |
| 01000 |  |  |  |  |  |  |  |  | 01000 |
| 01001 |  |  |  |  |  |  |  |  | 01001 |
| 01010 |  |  |  |  |  |  |  |  | 01010 |
| 01011 |  |  |  |  |  |  |  |  | 01011 |
| 01100 |  |  |  |  |  |  |  |  | 01100 |
| 01101 |  |  |  |  |  |  |  |  | 01101 |
| 01110 |  |  |  |  |  |  |  |  | 01110 |
| 01111 |  |  |  |  |  |  |  |  | 01111 |
| 10000 |  |  |  |  |  |  |  |  | 10000 |
| 10001 |  |  |  |  |  |  |  |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  |  |  | 10011 |
| 10100 |  |  |  |  |  |  |  |  | 10100 |
| 10101 |  |  |  |  |  |  |  |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  | 10110 |
| 10111 |  |  |  |  |  |  |  |  | 1011 |
| 11000 |  |  |  |  |  |  |  |  | 11000 |
| 11001 |  |  |  |  |  |  |  |  | 11001 |
| 11010 |  |  |  |  |  |  |  |  | 11010 |
| 11011 |  |  |  |  |  |  |  |  | 11011 |
| 11100 |  |  |  |  |  |  |  |  | 11100 |
| 11101 |  |  |  |  |  |  |  |  | 11101 |
| 11110 |  |  |  |  |  |  |  |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  | 11111 |
|  | 010000 | 010001 | 010010 | 010011 | 010100 | 010101 | 010110 | 010111 |  |

## Version 3.0

Table 15:EXT04: Extended Opcode Map for Primary Opcode 4 (opcode bits 21:31) (Sheet 4 of 8 )

|  | 011000 | 011001 | 011010 | 011011 | 011100 | 011101 | 011110 | 011111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 |  |  |  |  |  |  |  |  | 00000 |
| 00001 |  |  |  |  |  |  |  |  | 00001 |
| 00010 |  |  |  |  |  |  |  |  | 00010 |
| 00011 |  |  |  |  |  |  |  |  | 00011 |
| 00100 |  |  |  |  |  |  |  |  | 00100 |
| 00101 |  |  |  |  |  |  |  |  | 00101 |
| 00110 |  |  |  |  |  |  |  |  | 00110 |
| 00111 |  |  |  |  |  |  |  |  | 00111 |
| 01000 |  |  |  |  |  |  |  |  | 01000 |
| 01001 |  |  |  |  |  |  |  |  | 01001 |
| 01010 |  |  |  |  |  |  |  |  | 01010 |
| 01011 |  |  |  |  |  |  |  |  | 01011 |
| 01100 |  |  |  |  |  |  |  |  | 01100 |
| 01101 |  |  |  |  |  |  |  |  | 01101 |
| 01110 |  |  |  |  |  |  |  |  | 01110 |
| 01111 |  |  |  |  |  |  |  |  | 01111 |
| 10000 |  |  |  |  |  |  |  |  | 10000 |
| 10001 |  |  |  |  |  |  |  |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  |  |  | 10011 |
| 10100 |  |  |  |  |  |  |  |  | 10100 |
| 10101 |  |  |  |  |  |  |  |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  | 10110 |
| 10111 |  |  |  |  |  |  |  |  | 10111 |
| 11000 |  |  |  |  |  |  |  |  | 11000 |
| 11001 |  |  |  |  |  |  |  |  | 11001 |
| 11010 |  |  |  |  |  |  |  |  | 11010 |
| 11011 |  |  |  |  |  |  |  |  | 11011 |
| 11100 |  |  |  |  |  |  |  |  | 11100 |
| 11101 |  |  |  |  |  |  |  |  | 11101 |
| 11110 |  |  |  |  |  |  |  |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  | 11111 |
|  | 011000 | 011001 | 011010 | 011011 | 011100 | 011101 | 011110 | 011111 |  |

Table 15:EXT04: Extended Opcode Map for Primary Opcode 4 (opcode bits 21:31) (Sheet 5 of 8)


## Version 3.0

Table 15:EXT04: Extended Opcode Map for Primary Opcode 4 (opcode bits 21:31) (Sheet 6 of 8 )


Table 15:EXT04: Extended Opcode Map for Primary Opcode 4 (opcode bits 21:31) (Sheet 7 of 8 )


Version 3.0

Table 15:EXT04: Extended Opcode Map for Primary Opcode 4 (opcode bits 21:31) (Sheet 8 of 8 )


Table 16:XPND04-1: Expanded Opcode Map for $\mathrm{PO}=4 \mathrm{XO}=0 \mathrm{~b} 10110$ _ 000001 (opcode bits 21:30)

|  | 000 | 001 | 010 | 011 |
| :---: | :---: | :---: | :---: | :---: |
| 00 | $\left\|\begin{array}{cr} 00000 & \\ & \text { bcdctsq. } \\ \text { v3.0 } & \\ & \end{array}\right\|$ |  | $\left\|\begin{array}{rr} 00010 & \\ { }_{\mathrm{v} 3.0} & \text { bcdcfsq. } \\ & \mathrm{vx} \end{array}\right\|$ |  |
| 01 |  |  |  |  |
| 10 |  |  |  |  |
| 11 |  |  |  |  |
|  | 000 | 001 | 010 | 011 |



Table 17:XPND04-2: Expanded Opcode Map for $\mathrm{PO}=4 \mathrm{XO}=0 \mathrm{~b} 11110$ _ 000001 (opcode bits 21:30)

|  | 000 | 001 | 010 | 011 |
| :---: | :---: | :---: | :---: | :---: |
| 00 | $\left[\begin{array}{l}00000 \\ \text { binvalid\} }\end{array}\right]$ |  | $\begin{array}{\|cr\|} \hline 00{ }^{010} \text { bcdcfsq. } & \mathrm{I} \\ \mathrm{v}^{\mathrm{I} .0} & \mathrm{vx} \\ \hline \end{array}$ |  |
| 01 |  |  |  |  |
| 10 |  |  |  |  |
| 11 |  |  |  |  |
|  | 000 | 001 | 010 | 011 |




Table 18:XPND04-3: Expanded Opcode Map for $\mathrm{PO}=4 \mathrm{XO}=0 \mathrm{~b} 11000$ _ 000010 (opcode bits 21:30)


Version 3.0

Table 19:EXT19: Extended Opcode Map for Primary Opcode 19 (opcode bits 21:30) (Sheet 1 of 4)


Table 19:EXT19: Extended Opcode Map for Primary Opcode 19 (opcode bits 21:30) (Sheet 2 of 4)

|  | 01000 | 01001 | 01010 | 01011 | 01100 | 01101 | 01110 | 01111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 |  |  |  |  |  |  |  |  | 00000 |
| 00001 |  |  |  |  |  |  |  |  | 00001 |
| 00010 |  |  |  |  |  |  |  |  | 00010 |
| 00011 |  |  |  |  |  |  |  |  | 00011 |
| 00100 |  |  |  |  |  |  |  |  | 00100 |
| 00101 |  |  |  |  |  |  |  |  | 00101 |
| 00110 |  |  |  |  |  |  |  |  | 00110 |
| 00111 |  |  |  |  |  |  |  |  | 00111 |
| 01000 |  |  |  |  |  |  |  |  | 01000 |
| 01001 |  |  |  |  |  |  |  |  | 01001 |
| 01010 |  |  |  |  |  |  |  |  | 01010 |
| 01011 |  |  |  |  |  |  |  |  | 01011 |
| 01100 |  |  |  |  |  |  |  |  | 01100 |
| 01101 |  |  |  |  |  |  |  |  | 01101 |
| 01110 |  |  |  |  |  |  |  |  | 01110 |
| 01111 |  |  |  |  |  |  |  |  | 01111 |
| 10000 |  |  |  |  |  |  |  |  | 10000 |
| 10001 |  |  |  |  |  |  |  |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  |  |  | 10011 |
| 10100 |  |  |  |  |  |  |  |  | 10100 |
| 10101 |  |  |  |  |  |  |  |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  | 10110 |
| 10111 |  |  |  |  |  |  |  |  | 10111 |
| 11000 |  |  |  |  |  |  |  |  | 11000 |
| 11001 |  |  |  |  |  |  |  |  | 11001 |
| 11010 |  |  |  |  |  |  |  |  | 11010 |
| 11011 |  |  |  |  |  |  |  |  | 11011 |
| 11100 |  |  |  |  |  |  |  |  | 11100 |
| 11101 |  |  |  |  |  |  |  |  | 11101 |
| 11110 |  |  |  |  |  |  |  |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  | 11111 |
|  | 01000 | 01001 | 01010 | 01011 | 01100 | 01101 | 01110 | 01111 |  |

## Version 3.0

Table 19:EXT19: Extended Opcode Map for Primary Opcode 19 (opcode bits 21:30) (Sheet 3 of 4)


| 10100 | 10101 | 10110 | 10111 |  |
| :---: | :---: | :---: | :---: | :---: |
|  |  |  |  | 00000 |
|  |  |  |  | 00001 |
|  |  |  |  | 00010 |
|  |  |  |  | 00011 |
|  |  | 00100 10110 II <br> P1 isync  <br>  XL  |  | 00100 |
|  |  |  |  | 00101 |
|  |  |  |  | 00110 |
|  |  |  |  | 00111 |




| 10000 |
| :--- |
| 10001 |
| 10010 |
| 10011 |
| 10100 |
| 10101 |
| 10110 |
| 10111 |




Table 19:EXT19: Extended Opcode Map for Primary Opcode 19 (opcode bits 21:30) (Sheet 4 of 4)

|  | 11000 | 11001 | 11010 | 11011 | 11100 | 11101 | 11110 | 11111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 |  |  |  |  |  |  |  |  | 00000 |
| 00001 |  |  |  |  |  |  |  |  | 00001 |
| 00010 |  |  |  |  |  |  |  |  | 00010 |
| 00011 |  |  |  |  |  |  |  |  | 00011 |
| 00100 |  |  |  |  |  |  |  |  | 00100 |
| 00101 |  |  |  |  |  |  |  |  | 00101 |
| 00110 |  |  |  |  |  |  |  |  | 00110 |
| 00111 |  |  |  |  |  |  |  |  | 00111 |
| 01000 |  |  |  |  |  |  |  |  | 01000 |
| 01001 |  |  |  |  |  |  |  |  | 01001 |
| 01010 |  |  |  |  |  |  |  |  | 01010 |
| 01011 |  |  |  |  |  |  |  |  | 01011 |
| 01100 |  |  |  |  |  |  |  |  | 01100 |
| 01101 |  |  |  |  |  |  |  |  | 01101 |
| 01110 |  |  |  |  |  |  |  |  | 01110 |
| 01111 |  |  |  |  |  |  |  |  | 01111 |
| 10000 |  |  |  |  |  |  |  |  | 10000 |
| 10001 |  |  |  |  |  |  |  |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  |  |  | 10011 |
| 10100 |  |  |  |  |  |  |  |  | 10100 |
| 10101 |  |  |  |  |  |  |  |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  | 10110 |
| 10111 |  |  |  |  |  |  |  |  | 10111 |
| 11000 |  |  |  |  |  |  |  |  | 11000 |
| 11001 |  |  |  |  |  |  |  |  | 11001 |
| 11010 |  |  |  |  |  |  |  |  | 11010 |
| 11011 |  |  |  |  |  |  |  |  | 11011 |
| 11100 |  |  |  |  |  |  |  |  | 11100 |
| 11101 |  |  |  |  |  |  |  |  | 11101 |
| 11110 |  |  |  |  |  |  |  |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  | 11111 |
|  | 11000 | 11001 | 11010 | 11011 | 11100 | 11101 | 11110 | 11111 |  |

## Version 3.0

Table 20:EXT31: Extended Opcode Map for Primary Opcode 31 (opcode bits 21:30) (Sheet 1 of 4)

|  | 00000 | 00001 | 00010 | 00011 |
| :---: | :---: | :---: | :---: | :---: |
| 00000 | $\|$00000 00000 I <br> P1 $\mathbf{c m p}$ X |  |  |  |
| 00001 | $\left\|\begin{array}{lll} 00001 & 00000 & \mathrm{I} \\ \text { P1 } & \mathbf{c m p l} & \mathrm{X} \end{array}\right\|$ |  |  |  |
| 00010 | \{reserved\} |  |  |  |
| 00011 |  |  |  |  |
| 00100 | $\begin{array}{\|ccc\|} \hline 00100 & 0000 \\ & \text { setb } & \\ \text { v3.0 } & & x \\ \hline \end{array}$ |  |  | \{reserved\} |
| 00101 |  |  |  | \{reserved\} |
| 00110 |   <br> 00110 00000 <br> v3.0 cmprb <br>  $x$ |  |  |  |
| 00111 | $\left\|\begin{array}{ccc} 00111 & 00000 & \mathrm{I} \\ \text { v3.0 } & \text { cmpeqb } & \mathrm{x} \end{array}\right\|$ |  |  |  |


| 00100 | 00101 | 00110 |  | 00111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 00100 It <br> P1 tw  |  | $\begin{array}{ll} 00000 & 00110 \\ \text { v2.03 } & \text { IvsI } \end{array}$ |  |  | 1 <br> $\times$ |
| \{reserved\} |  | $\begin{array}{\|l\|l\|} \hline 00001 & 00110 \\ \text { v2.03 } & \text { Ivsr } \\ \hline \end{array}$ |  | $\begin{aligned} & 0000100111 \\ & \\ & \text { v2.03 } \\ & \text { Ivehx } \end{aligned}$ | I |
| $\begin{array}{\|lll} \hline 00010 & 00100 & \text { Id } \\ \text { PPC } & \text { td } & \\ \hline \end{array}$ |  |  |  | $\begin{array}{\|cc\|} \hline 00010 & 00111 \\ \text { Ivewx } \\ \text { v2.03 } \end{array}$ | I $\times$ $\times$ |
| \{reserved\} |  |  |  | $\begin{array}{\|l\|l\|} \hline 00011 & 00111 \\ \hline \text { v2.03 } & \text { Ivx } \\ \hline \end{array}$ |  |
|  |  |  |  | $\begin{array}{\|cc} 00100 & 00111 \\ \text { v2.03 } & \text { stvebx } \\ \hline \end{array}$ | I |
| \{reserved\} |  |  |  | 00101 <br>  <br> v2.03 <br> velvehx | I |
|  |  |  |  | $\left\lvert\, \begin{array}{cc} 00110 & 00111 \\ \text { v2.03 } & \text { stvewx } \end{array}\right.$ | I |
| \{reserved\} |  |  |  | $\begin{array}{\|l\|l\|} \hline 00111 & 00111 \\ \hline \text { v2.03 } & \text { stvx } \\ \hline \end{array}$ | I |


|  |
| :--- |
| 00000 |
| 00001 |
| 00010 |
| 00011 |
| 00100 |
| 00101 |
| 00110 |
| 00111 |



|  | \{reserved\} |  | 01000 |
| :---: | :---: | :---: | :---: |
| \{reserved\} |  |  | 01001 |
|  |  |  | 01010 |
| \{reserved\} |  | $\begin{array}{\|lll\|} \hline 01011 & 00111 & \text { I } \\ \text { v2.03 } & \text { IvxI } & \\ \hline \end{array}$ | 01011 |
|  |  |  | 01100 |
| \{reserved\} |  |  | 01101 |
|  | \{reserved\} |  | 01110 |
| \{reserved\} | \{reserved\} | 01111 00111 <br> v2.03 stvxI | 01111 |




|




Table 20:EXT31: Extended Opcode Map for Primary Opcode 31 (opcode bits 21:30) (Sheet 2 of 4)


Version 3.0

Table 20:EXT31: Extended Opcode Map for Primary Opcode 31 (opcode bits 21:30) (Sheet 3 of 4)


Table 20:EXT31: Extended Opcode Map for Primary Opcode 31 (opcode bits 21:30) (Sheet 4 of 4)


## Version 3.0

Table 21:EXT59: Extended Opcode Map for Primary Opcode 59 (opcode bits 21:30) (Sheet 1 of 4)


Table 21:EXT59: Extended Opcode Map for Primary Opcode 59 (opcode bits 21:30) (Sheet 2 of 4)

|  | 01000 | 01001 | 01010 | 01011 | 01100 | 01101 | 01110 | 01111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 |  |  |  |  |  |  |  |  | 00000 |
| 00001 |  |  |  |  |  |  |  |  | 00001 |
| 00010 |  |  |  |  |  |  |  |  | 00010 |
| 00011 |  |  |  |  |  |  |  |  | 00011 |
| 00100 |  |  |  |  |  |  |  |  | 00100 |
| 00101 |  |  |  |  |  |  |  |  | 00101 |
| 00110 |  |  |  |  |  |  |  |  | 00110 |
| 00111 |  |  |  |  |  |  |  |  | 00111 |
| 01000 |  |  |  |  |  |  |  |  | 01000 |
| 01001 |  |  |  |  |  |  |  |  | 01001 |
| 01010 |  |  |  |  |  |  |  |  | 01010 |
| 01011 |  |  |  |  |  |  |  |  | 01011 |
| 01100 |  |  |  |  |  |  |  |  | 01100 |
| 01101 |  |  |  |  |  |  |  |  | 01101 |
| 01110 |  |  |  |  |  |  |  |  | 01110 |
| 01111 |  |  |  |  |  |  |  |  | 01111 |
| 10000 |  |  |  |  |  |  |  |  | 10000 |
| 10001 |  |  |  |  |  |  |  |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  |  |  | 10011 |
| 10100 |  |  |  |  |  |  |  |  | 10100 |
| 10101 |  |  |  |  |  |  |  |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  | 10110 |
| 10111 |  |  |  |  |  |  |  |  | 1011 |
| 11000 |  |  |  |  |  |  |  |  | 11000 |
| 11001 |  |  |  |  |  |  |  |  | 11001 |
| 11010 |  |  |  |  |  |  | $\begin{array}{\|l\|l\|l\|} \hline 11010 & \begin{array}{c} 11110 \\ \text { vefids[.06 } \end{array} \\ \hline \end{array}$ |  | 11010 |
| 11011 |  |  |  |  |  |  |  |  | 11011 |
| 11100 |  |  |  |  |  |  |  |  | 11100 |
| 11101 |  |  |  |  |  |  |  |  | 11101 |
| 11110 |  |  |  |  |  |  | $\begin{aligned} & 111110 \text { 01110 } \\ & \text { fcfidus[.] } \\ & \text { v2.06 } \end{aligned}$ |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  | 11111 |
|  | 01000 | 01001 | 01010 | 01011 | 01100 | 01101 | 01110 | 01111 |  |

## Version 3.0

Table 21:EXT59: Extended Opcode Map for Primary Opcode 59 (opcode bits 21:30) (Sheet 3 of 4)

|  | 10000 | 10001 | 10010 | 10011 | 10100 | 10101 | 10110 | 10111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 |  |  |  |  | $\begin{array}{\|c\|c\|} \hline / 1 / / / 10100 \\ \text { PPC } & \text { fsubs[.] } \\ \hline \end{array}$ | $\begin{aligned} & \text { I } \mid / / / / / 10101 \\ & \text { A Padds[.] } \end{aligned}$ | I $\left\lvert\, / / / / / \frac{10110}{\text { fsqrts[.] }}\right.$ |  | 00000 |
| 00001 |  |  |  |  | $\|$TIIII 10100 <br> fsubs[.] <br> \{invalid\} | $\|$IIII/ 10101 <br> fadds[.] <br> finvalid\} | $\|$/IIII/ 10110 <br> fsqrts[.] <br> finvalid\} |  | 00001 |
| 00010 |  |  |  |  |  |  |  |  | 00010 |
| 00011 |  |  |  |  |  |  |  |  | 00011 |
| 00100 |  |  |  |  |  |  |  |  | 00100 |
| 00101 |  |  |  |  |  |  |  |  | 00101 |
| 00110 |  |  |  |  |  |  |  |  | 00110 |
| 0011 |  |  |  |  |  |  |  |  | 0011 |
| 01000 |  |  |  |  |  |  |  |  | 01000 |
| 01001 |  |  |  |  |  |  |  |  | 01001 |
| 01010 |  |  |  |  |  |  |  |  | 01010 |
| 01011 |  |  |  |  |  |  |  |  | 01011 |
| 01100 |  |  |  |  |  |  |  |  | 01100 |
| 01101 |  |  |  |  |  |  |  |  | 01101 |
| 01110 |  |  |  |  |  |  |  |  | 01110 |
| 01111 |  |  |  |  |  |  |  |  | 01111 |
| 10000 |  |  |  |  |  |  |  |  | 10000 |
| 10001 |  |  |  |  |  |  |  |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  |  |  | 10011 |
| 10100 |  |  |  |  |  |  |  |  | 10100 |
| 10101 |  |  |  |  |  |  |  |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  | 10110 |
| 10111 |  |  |  |  |  |  |  |  | 10111 |
| 11000 |  |  |  |  |  |  |  |  | 11000 |
| 11001 |  |  |  |  |  |  |  |  | 11001 |
| 11010 |  |  |  |  |  |  |  |  | 11010 |
| 11011 |  |  |  |  |  |  |  |  | 11011 |
| 11100 |  |  |  |  |  |  |  |  | 11100 |
| 11101 |  |  |  |  |  |  |  |  | 11101 |
| 11110 |  |  |  |  |  |  |  |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  | 11111 |
|  | 10000 | 10001 | 10010 | 10011 | 10100 | 10101 | 10110 | 10111 |  |

Table 21:EXT59: Extended Opcode Map for Primary Opcode 59 (opcode bits 21:30) (Sheet 4 of 4)

|  | 11000 | 11001 | 11010 | 11011 | 11100 | 11101 | 11110 | 11111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 | $\begin{array}{\|cc\|} \hline 1 / 1 / / 11000 \\ \text { PPC } & \text { fres[.] } \end{array}$ | $\begin{gathered} \text { I } \\ \text { A PPC } \end{gathered} \begin{gathered} 11001 \\ \text { fmuls[.] } \end{gathered}$ | $\begin{aligned} & \text { I///I// 11010 } \\ & \text { A } \mid \text { v2.02 } \\ & \text { frsqres[.] } \end{aligned}$ |  |  |  | $\begin{aligned} & \text { I } \ldots . . .11110 \\ & \text { A PPC } \end{aligned}$ | $\begin{aligned} & \text { I } \cdot \cdots \cdot 11111 \\ & \text { APPC } \end{aligned}$ | 00000 |
| 00001 | fres[.] |  | frsqrtes[.] |  |  |  |  |  | 00001 |
| 00010 |  |  |  |  |  |  |  |  | 00010 |
| 00011 |  |  |  |  |  |  |  |  | 00011 |
| 00100 |  |  |  |  |  |  |  |  | 00100 |
| 00101 |  |  |  |  |  |  |  |  | 00101 |
| 00110 |  |  |  |  |  |  |  |  | 00110 |
| 00111 |  |  |  |  |  |  |  |  | 00111 |
| 01000 |  |  |  |  |  |  |  |  | 01000 |
| 01001 |  |  |  |  |  |  |  |  | 01001 |
| 01010 |  |  |  |  |  |  |  |  | 01010 |
| 01011 |  |  |  |  |  |  |  |  | 01011 |
| 01100 |  |  |  |  |  |  |  |  | 01100 |
| 01101 |  |  |  |  |  |  |  |  | 01101 |
| 01110 |  |  |  |  |  |  |  |  | 01110 |
| 01111 |  |  |  |  |  |  |  |  | 01111 |
| 10000 |  |  |  |  |  |  |  |  | 10000 |
| 10001 |  |  |  |  |  |  |  |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  |  |  | 10011 |
| 10100 |  |  |  |  |  |  |  |  | 10100 |
| 10101 |  |  |  |  |  |  |  |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  | 10110 |
| 10111 |  |  |  |  |  |  |  |  | 1011 |
| 11000 |  |  |  |  |  |  |  |  | 11000 |
| 11001 |  |  |  |  |  |  |  |  | 11001 |
| 11010 |  |  |  |  |  |  |  |  | 11010 |
| 11011 |  |  |  |  |  |  |  |  | 11011 |
| 11100 |  |  |  |  |  |  |  |  | 11100 |
| 11101 |  |  |  |  |  |  |  |  | 11101 |
| 11110 |  |  |  |  |  |  |  |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  | 11111 |
|  | 11000 | 11001 | 11010 | 11011 | 11100 | 11101 | 11110 | 11111 |  |

## Version 3.0

Table 22:EXT60: Extended Opcode Map for Primary Opcode 60 (opcode bits 21:30) (Sheet 1 of 4)

|  | 00000 | 00001 | 00010 | 00011 | 00100 | 00101 | 00110 | 00111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 | $\left\lvert\, \begin{gathered} 00000 \text { 000.. } \\ \text { xsaddsp } \end{gathered}\right.$ |  |  |  |  |  |  |  | 00000 |
| 00001 | $\begin{aligned} & 00001 \text { 000.. } \\ & \text { xssubsp } \\ & \text { v2.07 } \end{aligned}$ |  |  |  | $\qquad$ |  |  |  | 00001 |
| 00010 | $\left\lvert\, \begin{array}{ll} 00010 & 000 . . \\ \text { xsmulsp } \end{array}\right.$ |  |  |  | 00010 001.. xsmsubasp v2.07 XX |  |  |  | 00010 |
| 00011 | $\begin{array}{\|cc} \hline 00011 & 000 . . \\ & \text { xsdivsp } \\ \text { v2.07 } \end{array}$ |  |  |  |  |  |  |  | 00011 |
| 00100 | $\left\lvert\, \begin{aligned} & 00100 \text { 000.. } \\ & \text { x2.06 } \\ & \text { vadddp } \end{aligned}\right.$ |  |  |  | $\begin{gathered} 00100 \text { 001.. } \\ \text { xsmaddadp } \\ \text { vx3 } \end{gathered}$ |  |  |  | 00100 |
| 00101 | $\begin{aligned} & 00101 \text { 000.. } \\ & \text { v2.06 } \end{aligned}$ |  |  |  |  |  |  |  | 00101 |
| 00110 | $\begin{gathered} 00110 \text { 000.. } \\ \text { xsmuldp } \\ \text { v2.06 } \end{gathered}$ |  |  |  | $\begin{aligned} & \text { xsmsubadp } \\ & \text { v2.06 } \\ & \text { xX3 } \\ & \hline \end{aligned}$ |  |  |  | 00110 |
| 00111 | $\begin{array}{\|cc} \hline 00111 & 000 . . \\ \\ \text { v2.06 } & \\ \hline \text { xsdivdp } \\ \hline \end{array}$ |  |  |  | $\qquad$ |  |  |  | 00111 |
| 01000 | $\left.\right\|_{\text {v2.06 }} ^{01000 \text { 000.. }} \text { xvaddsp }$ |  |  |  | $\qquad$ |  |  |  | 01000 |
| 01001 |  |  |  |  | 01001 001..  <br> xvmaddmsp  <br> v2.06 XX3 |  |  |  | 01001 |
| 01010 | $\begin{array}{\|l} \text { 01010 000.. } \\ \text { v2.06 } \end{array}$ |  |  |  |  |  |  |  | 01010 |
| 01011 | $\begin{array}{\|cc} \hline 01011 & \text { 000.. } \\ \text { vvdivsp } \\ \text { v2.06 } \end{array}$ |  |  |  | $\qquad$ |  |  |  | 01011 |
| 01100 | $\begin{aligned} & \text { 01100 000.. } \\ & \text { v2.06 } x \text { xvadddp } \\ & \hline \end{aligned}$ |  |  |  |  |  |  |  | 01100 |
| 01101 | $\begin{array}{\|l} 01101 \text { 000.. } \\ \text { v2.06 } \\ \hline \end{array}$ |  |  |  | $\begin{aligned} & \begin{array}{l} 01101 \text { 001.. } \\ \text { xvmaddmdp } \\ \text { v2.06 } \end{array} \quad \text { XX3 } \end{aligned}$ |  |  |  | 01101 |
| 01110 | $\begin{aligned} & \text { 01110 000.. } \\ & \text { xvmuldp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  | $\qquad$ |  |  |  | 01110 |
| 01111 | $\begin{aligned} & 01111 \text { 000.. } \\ & \text { xvdivdp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  |  |  |  |  | 01111 |
| 10000 | $\begin{aligned} & 10000 \text { 000.. } \\ & \text { vsmaxcdp } \\ & \text { v3.0 } \end{aligned}$ |  |  |  | 10000 001..  <br> xsnmaddasp  <br> V2.07 XX3 |  |  |  | 10000 |
| 10001 | $\begin{aligned} & 10001 \text { 000... } \\ & \text { xsmincdp } \\ & \text { v3.0 } \end{aligned}$ |  |  |  | 10001 001.. <br> xsnmaddmsp <br> v2.07 |  |  |  | 10001 |
| 10010 | $\begin{aligned} & 10010 \quad 000 . . \\ & \begin{array}{l} \text { vs.0 } \\ \text { vsmaxjdp } \end{array} \\ & \hline \end{aligned}$ |  |  |  | $\begin{gathered} 10010 \text { 001.. } \\ \text { xsnmsubasp } \\ \mathrm{V} 2.07 \end{gathered}$ |  |  |  | 10010 |
| 10011 | $\begin{aligned} & 10011 \text { 000.. } \\ & \text { xsminjdp } \\ & \text { v3.0 } \end{aligned}$ |  |  |  | 10011 001.. <br> xsnmsubmsp <br> v2.07 |  |  |  | 10011 |
| 10100 | $\begin{aligned} & 10100 \text { 000.. } \\ & \text { xsmaxdp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  | $\begin{gathered} 10100 \text { 001.. } \\ \text { xsnmaddadp } \\ \text { v2.06 } \quad \text { XX3 } \end{gathered}$ |  |  |  | 10100 |
| 10101 | $\begin{aligned} & 10101 \text { 000... } \\ & \text { xsmindp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  | $\begin{gathered} 10101 \text { 001.. } \\ \begin{array}{l} \text { xsnmaddmdp } \\ \text { v2.06 } \end{array} \quad \text { XX3 } \end{gathered}$ |  |  |  | 10101 |
| 10110 | $\begin{aligned} & 10110 \text { 000... } \\ & \quad \text { xscpsgndp } \\ & \text { l2.06 } \end{aligned}$ |  |  |  | $\begin{aligned} & 10110 \text { 001.. } \\ & \text { xsnmsubadp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  | 10110 |
| 1011 |  |  |  |  | $\begin{aligned} & 10111 \text { 001.. } \\ & \text { xsnmsubmdp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  | 10111 |
| 11000 | $\begin{aligned} & 11000 \text { 000.. } \\ & \text { xvmaxsp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  | $\begin{aligned} & 11000 \text { 001.. } \\ & \text { xvnmaddasp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  | 11000 |
| 11001 | $\begin{array}{\|l} 11001 \text { 000... } \\ \\ \text { vvminsp } \end{array}$ |  |  |  | $\begin{aligned} & 11001 \text { 001.. } \\ & \begin{array}{l} \text { xvnmaddmsp } \\ \text { v2.06 } \end{array} \quad \text { XX3 } \end{aligned}$ |  |  |  | 11001 |
| 11010 | $\begin{aligned} & 11010 \text { 000.. } \\ & \text { xvcpsgnsp } \\ & \mathrm{v} 2.06 \end{aligned}$ |  |  |  | $\begin{aligned} & 11010 \text { 001.. } \\ & \text { xvnmsubasp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  | 11010 |
| 11011 | $\begin{array}{\|l\|l} \hline 11011 \begin{array}{l} \text { 000.. } \\ \text { xviexpsp } \\ \text { v3.0 } \end{array} \\ \hline \end{array}$ |  |  |  |  |  |  |  | 11011 |
| 11100 | $\begin{aligned} & \text { 11100 000.. } \\ & \text { v2.06 } \\ & \hline \end{aligned}$ |  |  |  | $\begin{aligned} & 11100 \text { 001.. } \\ & \begin{array}{l} \text { xvnmaddadp } \\ \text { v2.06 } \end{array} \quad \mathrm{XX3} \end{aligned}$ |  |  |  | 11100 |
| 11101 | $\begin{aligned} & 11101 \text { 000.. } \\ & \text { xvmindp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  |  |  |  |  | 11101 |
| 11110 | $\begin{aligned} & 11110 \text { 000.. } \\ & \text { xvcpsgndp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  |  |  |  |  | 11110 |
| 11111 |  |  |  |  | $\begin{aligned} & 11111 \text { 001.. } \\ & \text { xvnmsubmdp } \\ & \text { v2.06 } \end{aligned}$ |  |  |  | 11111 |
|  | 00000 | 00001 | 00010 | 00011 | 00100 | 00101 | 00110 | 00111 |  |

Table 22:EXT60: Extended Opcode Map for Primary Opcode 60 (opcode bits 21:30) (Sheet 2 of 4)


Version 3.0

Table 22:EXT60: Extended Opcode Map for Primary Opcode 60 (opcode bits 21:30) (Sheet 3 of 4)


Table 22:EXT60: Extended Opcode Map for Primary Opcode 60 (opcode bits 21:30) (Sheet 4 of 4)

|  | 11000 |  | 11001 | 11010 | 11011 | 11100 | 11101 | 11110 | 11111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 | $\left\lvert\, \begin{array}{\|cc} \ldots & 11 . . \\ \text { v2.06 } & \text { xxsel } \end{array}\right.$ |  |  |  |  |  |  |  |  | 00000 |
| 00001 |  |  |  |  |  |  |  |  |  | 00001 |
| 00010 |  |  |  |  |  |  |  |  |  | 00010 |
| 00011 |  |  |  |  |  |  |  |  |  | 00011 |
| 00100 |  |  |  |  |  |  |  |  |  | 00100 |
| 00101 |  |  |  |  |  |  |  |  |  | 00101 |
| 00110 |  |  |  |  |  |  |  |  |  | 00110 |
| 00111 |  |  |  |  |  |  |  |  |  | 00111 |
| 01000 |  |  |  |  |  |  |  |  |  | 01000 |
| 01001 |  |  |  |  |  |  |  |  |  | 01001 |
| 01010 |  |  |  |  |  |  |  |  |  | 01010 |
| 01011 |  |  |  |  |  |  |  |  |  | 01011 |
| 01100 |  |  |  |  |  |  |  |  |  | 01100 |
| 01101 |  |  |  |  |  |  |  |  |  | 01101 |
| 01110 |  |  |  |  |  |  |  |  |  | 01110 |
| 01111 |  |  |  |  |  |  |  |  |  | 01111 |
| 10000 |  |  |  |  |  |  |  |  |  | 10000 |
| 10001 |  |  |  |  |  |  |  |  |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  |  |  |  | 10011 |
| 10100 |  |  |  |  |  |  |  |  |  | 10100 |
| 10101 |  |  |  |  |  |  |  |  |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  |  | 10110 |
| 10111 |  |  |  |  |  |  |  |  |  | 10111 |
| 11000 |  |  |  |  |  |  |  |  |  | 11000 |
| 11001 |  |  |  |  |  |  |  |  |  | 11001 |
| 11010 |  |  |  |  |  |  |  |  |  | 11010 |
| 11011 |  |  |  |  |  |  |  |  |  | 11011 |
| 11100 |  |  |  |  |  |  |  |  |  | 11100 |
| 11101 |  |  |  |  |  |  |  |  |  | 11101 |
| 11110 |  |  |  |  |  |  |  |  |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  |  | 11111 |
|  | 11000 |  | 11001 | 11010 | 11011 | 11100 | 11101 | 11110 | 11111 |  |

Table 23:XPND60-1: Expanded Opcode Map for $\mathrm{PO}=60 \mathrm{XO}=0 \mathrm{~b} 01011$ _ 01000 (opcode bits 21:30)


Table 24:XPND60-2: Expanded Opcode Map for PO=60 XO=0b10101_1011. (opcode bits 21:30)


Table 25:XPND60-3: Expanded Opcode Map for $\mathrm{PO}=60 \mathrm{XO}=0 \mathrm{~b} 11101$ _ 1011 . (opcode bits 21:30)


Table 26:EXT63: Extended Opcode Map for Primary Opcode 63 (opcode bits 21:30) (Sheet 1 of 4)


## Version 3.0

Table 26:EXT63: Extended Opcode Map for Primary Opcode 63 (opcode bits 21:30) (Sheet 2 of 4)

|  | 01000 | 01001 | 01010 | 01011 | 01100 | 01101 | 01110 | 01111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 |  |  |  |  | 00000 01100 <br> P1 frsp[.] |  |  | $\begin{array}{\|c\|c\|} \hline 100000 & 01111 \\ \text { x } & \text { fatiwz[.] } \\ & \\ \hline \end{array}$ | 00000 |
| 00001 | P100001 01000 <br> fneg[.]  |  |  |  |  |  |  |  | 00001 |
| 00010 | $\left[\begin{array}{cc} 20010 & 01000 \\ P 1 & f m r[.] \end{array}\right.$ |  |  |  |  |  |  |  | 00010 |
| 00011 |  |  |  |  |  |  |  |  | 00011 |
| 00100 | $\begin{array}{\|c\|c} \hline 00100 & 01000 \\ \text { P1 } & \text { fnabs[.] } \\ \hline \end{array}$ |  |  |  |  |  | $\begin{array}{\|c\|c\|} \hline 00100 & 01110 \\ \text { v2.06 } \end{array}$ |  | 00100 |
| 00101 |  |  |  |  |  |  |  |  | 00101 |
| 00110 |  |  |  |  |  |  |  |  | 00110 |
| 00111 |  |  |  |  |  |  |  |  | 00111 |
| 01000 | $\begin{array}{ll} \hline 1000001000 \\ \text { P1 } & \text { fabs[.] } \\ \hline \end{array}$ |  |  |  |  |  |  |  | 01000 |
| 01001 |  |  |  |  |  |  |  |  | 01001 |
| 01010 |  |  |  |  |  |  |  |  | 01010 |
| 01011 |  |  |  |  |  |  |  |  | 01011 |
| 01100 | $\begin{aligned} & \hline 01100 \\ & \mathrm{c} 2.02000 \\ & \mathrm{v} 2000 \\ & \text { frin[.] } \end{aligned}$ |  |  |  |  |  |  |  | 01100 |
| 01101 |  |  |  |  |  |  |  |  | 01101 |
| 01110 | $\begin{gathered} 01110 \\ \text { v20000 } \\ \text { v2000 } \\ \text { frip[.] } \end{gathered}$ |  |  |  |  |  |  |  | 01110 |
| 01111 |  |  |  |  |  |  |  |  | 01111 |
| 10000 |  |  |  |  |  |  |  |  | 10000 |
| 10001 |  |  |  |  |  |  |  |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  |  |  | 10011 |
| 10100 |  |  |  |  |  |  |  |  | 10100 |
| 10101 |  |  |  |  |  |  |  |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  | 10110 |
| 10111 |  |  |  |  |  |  |  |  | 10111 |
| 11000 |  |  |  |  |  |  |  |  | 11000 |
| 11001 |  |  |  |  |  |  | $\begin{array}{\|l\|l\|} \hline 11001 & 01110 \\ \text { PPC } & \text { fctid[.] } \\ \hline \end{array}$ | $\begin{aligned} & \text { I } \\ & \mathrm{x} \mid 1100101111 \\ & \text { PPC } \text { fctidz[] } \end{aligned}$ | 11001 |
| 11010 |  |  |  |  |  |  |  |  | 11010 |
| 11011 |  |  |  |  |  |  |  |  | 11011 |
| 11100 |  |  |  |  |  |  |  |  | 11100 |
| 11101 |  |  |  |  |  |  |  | $\left.\begin{array}{cc} \mathrm{I} \mid 11101 \text { e11111 } \\ \text { fatiduz[.] } & \mathrm{I} \\ \mathrm{x} 2.06 \end{array} \right\rvert\,$ | 11101 |
| 11110 |  |  |  |  |  |  | $\begin{aligned} & 111100 \text { 01110 } \\ & \text { fcfidu[.] } \\ & \text { v2.06 } \end{aligned}$ |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  | 11111 |
|  | 01000 | 01001 | 01010 | 01011 | 01100 | 01101 | 01110 | 01111 |  |

Table 26:EXT63: Extended Opcode Map for Primary Opcode 63 (opcode bits 21:30) (Sheet 3 of 4)

|  | 10000 | 10001 | 10010 | 10011 | 10100 | 10101 | 10110 | 10111 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 00000 |  |  |  |  | $\begin{array}{ll} \hline 1 / I / / 1 & 10100 \\ \text { Ps } & \\ \hline \end{array}$ |  |  | I $\ldots$ 10111 <br> A PPC  <br> fsel[ $[]$.   | 00000 |
| 00001 |  |  | $\overbrace{\{\text { invalid }\}}^{\text {flit/ } 10010} \mathrm{fdiv}[$.] |  | I/III/ 10100 <br> finvalid\} | /III/ 10101 finvalid\} fadd[.] |  |  | 00001 |
| 00010 |  |  |  |  |  |  |  |  | 00010 |
| 00011 |  |  |  |  |  |  |  |  | 00011 |
| 00100 |  |  |  |  |  |  |  |  | 00100 |
| 00101 |  |  |  |  |  |  |  |  | 00101 |
| 00110 |  |  |  |  |  |  |  |  | 00110 |
| 00111 |  |  |  |  |  |  |  |  | 00111 |
| 01000 |  |  |  |  |  |  |  |  | 01000 |
| 01001 |  |  |  |  |  |  |  |  | 01001 |
| 01010 |  |  |  |  |  |  |  |  | 01010 |
| 01011 |  |  |  |  |  |  |  |  | 01011 |
| 01100 |  |  |  |  |  |  |  |  | 01100 |
| 01101 |  |  |  |  |  |  |  |  | 01101 |
| 01110 |  |  |  |  |  |  |  |  | 01110 |
| 01111 |  |  |  |  |  |  |  |  | 01111 |
| 10000 |  |  |  |  |  |  |  |  | 10000 |
| 10001 |  |  |  |  |  |  |  |  | 10001 |
| 10010 |  |  |  |  |  |  |  |  | 10010 |
| 10011 |  |  |  |  |  |  |  |  | 10011 |
| 10100 |  |  |  |  |  |  |  |  | 10100 |
| 10101 |  |  |  |  |  |  |  |  | 10101 |
| 10110 |  |  |  |  |  |  |  |  | 10110 |
| 10111 |  |  |  |  |  |  |  |  | 1011 |
| 11000 |  |  |  |  |  |  |  |  | 11000 |
| 11001 |  |  |  |  |  |  |  |  | 11001 |
| 11010 |  |  |  |  |  |  |  |  | 11010 |
| 11011 |  |  |  |  |  |  |  |  | 11011 |
| 11100 |  |  |  |  |  |  |  |  | 11100 |
| 11101 |  |  |  |  |  |  |  |  | 11101 |
| 11110 |  |  |  |  |  |  |  |  | 11110 |
| 11111 |  |  |  |  |  |  |  |  | 11111 |
|  | 10000 | 10001 | 10010 | 10011 | 10100 | 10101 | 10110 | 10111 |  |

## Version 3.0

Table 26:EXT63: Extended Opcode Map for Primary Opcode 63 (opcode bits 21:30) (Sheet 4 of 4)


Table 27:XPND63-1: Expanded Opcode Map for $\mathrm{PO}=63$ XO=0b11001_ 00100 (opcode bits 21:30)

|  | 000 | 001 | 010 | 011 |
| :---: | :---: | :---: | :---: | :---: |
| 00 | $\left\|\begin{array}{ll} 00 \\ { }_{\text {v3.0 }}^{000} & \text { xsabsqp } \end{array}\right\|$ |  | $\|$$00{ }^{0010}$ I <br> vs.0  |  |
| 01 | $\begin{array}{\|ll\|} \hline 01000 & \mathrm{I} \\ \text { v3.0 } \end{array}$ |  |  |  |
| 10 | $\begin{array}{\|lr} \hline 10000 & \mathrm{I} \\ \mathrm{v}^{10.0} \begin{array}{c} \text { xsnegqp } \end{array} & \mathrm{x} \\ \hline \end{array}$ |  | $\begin{array}{\|ll} \hline 10{ }^{1010} & \text { vsxsigqp } \\ \text { v3.0 } \end{array}$ |  |
| 11 |  |  |  | ```\|\mp@code{l1 011 xssqrtqp[0] }``` |
|  | 000 | 001 | 010 | 011 |


| 100 | 101 | 110 | 111 |  |
| :---: | :---: | :---: | :---: | :---: |
|  |  |  |  | 00 |
|  |  |  |  | 01 |
|  |  |  |  | 10 |
|  |  |  |  | 11 |
| 100 | 101 | 110 | 111 |  |

Table 28:XPND63-2: Expanded Opcode Map for PO=63 XO=Ob11010_00100 (opcode bits 21:30)

|  | 000 | 001 | 010 | 011 |
| :---: | :---: | :---: | :---: | :---: |
| 00 |  | $\left\lvert\, \begin{array}{ll} 00001 \\ { }_{\text {v3.0 }}^{00} \mathbf{x s c v q p u w z ~} & \mathrm{I} \\ \hline \end{array}\right.$ |  |  |
| 01 |  | $\begin{array}{\|ll} \hline 01 \text { 001 } & \mathrm{I} \\ \text { vscvqpswz } & \mathrm{xscvq} \\ \hline \end{array}$ | $\left\lvert\, \begin{array}{ll} \text { 01 } \\ { }^{010} & \text { xscvsdqp } \\ \text { v3.0 } \end{array}\right.$ |  |
| 10 |  | 10001 <br> vscvqpudz <br> v3.0 <br> x |  |  |
| 11 |  | $\begin{array}{\|cc\|} \hline 11 \begin{array}{l} \text { 001 } \\ \text { vscvqpsdz } \\ \end{array} & \mathrm{x} \\ \hline \end{array}$ |  |  |
|  | 000 | 001 | 010 | 011 |


| 100 | 101 | 110 | 111 |
| :---: | :---: | :---: | :---: |
|  |  |  |  |
|  |  |  |  |
| $\begin{array}{\|ll\|} \hline 10 \begin{array}{ll} \text { 100 } \\ \text { xscvqpdp[o] } \\ \text { v3.0 } \end{array} & \mathrm{X} \\ \hline \end{array}$ |  |   <br> 1010 <br> vscudpqp <br> xscva $x$ |  |
|  |  |  |  |
| 100 | 101 | 110 | 111 |

## Appendix D. Power ISA Instruction Set Sorted by Opcode

This appendix lists all the instructions in the Power ISA, sorted by primary opcode, then by extended opcode bits 26:31 (if any), then by opcode bits 21:25 (if any), then by expanded opcode bits 11:15 (if any).


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 1 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 2 of 17)

| Instruction $^{\mathbf{1}}$ <br> 1 <br> $112345 \quad 67890 \quad 1211111112 \quad 22222 \quad 222233$ <br> 67890 <br> $12345 \quad 678901$ |  | $\begin{aligned} & \text { ro } \\ & \text { on } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & \underset{\sim}{0} \end{aligned}$ |  |  | + 0 0 0 0 0 0 | 0 <br> 0 <br> 0 <br> $\mathbf{0}$ | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $000100 \ldots . .11111 \ldots . .1111000011$ | VX | I | 348 | v2.07 |  |  | vpopentw | Vector Population Count Word |
| $000100 \ldots . .11111 \ldots . .11111000011$ | VX | I | 348 | v2.07 |  |  | vpopentd | Vector Population Count Dword |
| $000100 \ldots \ldots . . . . . . . .00000000100$ | VX | I | 318 | v2.03 |  |  | vrlb | Vector Rotate Left Byte |
|  | VX | I | 318 | v2.03 |  |  | vrlh | Vector Rotate Left Hword |
| $000100 \ldots \ldots . \ldots \ldots 000010000100$ | VX | I | 318 | v2.03 |  |  | vrlw | Vector Rotate Left Word |
| $000100 \ldots \ldots . \ldots \ldots . .100011000100$ | VX | 1 | 318 | v2.07 |  |  | vrld | Vector Rotate Left Dword |
| $000100 \ldots \ldots . \ldots \ldots 0.10000000100$ | VX | 1 | 319 | v2.03 |  |  | vslb | Vector Shift Left Byte |
| $000100 \ldots \ldots . \ldots \ldots, \ldots .00101000100$ | VX | I | 319 | v2.03 |  |  | vslh | Vector Shift Left Hword |
|  | VX | I | 319 | v2.03 |  |  | vslw | Vector Shift Left Word |
| $000100 \ldots \ldots . \ldots . . . . .100111000100$ | VX | I | 266 | v2.03 |  |  | vSI | Vector Shift Left |
| $000100 \ldots \ldots . \ldots \ldots 01000000100$ | VX | I | 320 | v2.03 |  |  | vsrb | Vector Shift Right Byte |
|  | VX | I | 320 | v2.03 |  |  | vsrh | Vector Shift Right Hword |
| $000100 \ldots \ldots . \ldots \ldots .101010000100$ | VX | I | 320 | v2.03 |  |  | vSTw | Vector Shift Right Word |
| $000100 \ldots \ldots . \ldots \ldots 0.1010 .1000100$ | VX | I | 266 | v2.03 |  |  | vSr | Vector Shift Right |
| 000100 $\ldots \ldots \ldots \ldots \ldots . . . . . . .1100000100$ | VX | , | 321 | v2.03 |  |  | vsrab | Vector Shift Right Algebraic Byte |
| $000100 \ldots \ldots . . . . . . . . .101101000100$ | VX | I | 321 | v2.03 |  |  | vsrah | Vector Shift Right Algebraic Hword |
| $000100 \ldots \ldots . \ldots \ldots 0.10 .110000100$ | VX | I | 321 | v2.03 |  |  | vsraw | Vector Shift Right Algebraic Word |
| $000100 \ldots \ldots . \ldots \ldots 0.1111000100$ | VX | 1 | 321 | v2.07 |  |  | vsrad | Vector Shift Right Algebraic Dword |
| 000100 $\ldots \ldots \ldots \ldots \ldots . . . .10000000100$ | VX | 1 | 315 | v2.03 |  |  | vand | Vector Logical AND |
| $000100 \ldots \ldots . \ldots \ldots 10001000100$ | VX | I | 315 | v2.03 |  |  | vandc | Vector Logical AND with Complement |
| $000100 \ldots \ldots \ldots \ldots 10010000100$ | VX | I | 316 | v2.03 |  |  | vor | Vector Logical OR |
| $000100 \ldots \ldots . \ldots \ldots .10011000100$ | VX | I | 316 | v2.03 |  |  | vxor | Vector Logical XOR |
| 000100 $\ldots \ldots \ldots \ldots \ldots . . . . .10100000100$ | VX | , | 316 | v2.03 |  |  | vnor | Vector Logical NOR |
| 000100 $\ldots \ldots \ldots \ldots \ldots .10101000100$ | VX | \| | 316 | v2.07 |  |  | vorc | Vector Logical OR with Complement |
| $000100 \ldots \ldots . . . .1 . . .10110000100$ | VX | 1 | 315 | v2.07 |  |  | vnand | Vector Logical NAND |
| $000100 \ldots \ldots . . . . . . . . .10111000100$ | VX | 1 | 319 | v2.07 |  |  | vsld | Vector Shift Left Dword |
| $000100 \ldots . .111111111111000000100$ | VX | I | 364 | v2.03 |  |  | mfvscr | Move From VSCR |
| $0001001111111111 \ldots . .11001000100$ | VX | I | 364 | v2.03 |  |  | mtvscr | Move To VSCR |
| $000100 \ldots \ldots . \ldots \ldots .11010000100$ | VX | I | 315 | v2.07 |  |  | veqv | Vector Logical Equivalence |
| $000100 \ldots \ldots . . . . .1 . . .11011000100$ | VX | 1 | 320 | v2.07 |  |  | VSrd | Vector Shift Right Dword |
| 000100 $\ldots \ldots \ldots \ldots \ldots .11100000100$ | X | I | 267 | v3.0 |  |  | vsiv | Vector Shift Right Variable |
| $000100 \ldots \ldots . \ldots \ldots \ldots 11101000100$ | X | 1 | 267 | v3.0 |  |  | vslv | Vector Shift Left Variable |
| $000100 \ldots \ldots . \ldots \ldots 0.10010000101$ | VX | 1 | 322 | v3.0 |  |  | vrlwmi | Vector Rotate Left Word then Mask Insert |
| $000100 \ldots \ldots . \ldots \ldots . . .100011000101$ | VX | I | 323 | v3.0 |  |  | vrldmi | Vector Rotate Left Dword then Mask Insert |
| $000100 \ldots \ldots . \ldots \ldots .00110000101$ | VX | I | 322 | v3.0 |  |  | vrlwnm | Vector Rotate Left Word then AND with Mask |
| $000100 \ldots \ldots . \ldots \ldots, \ldots 0111000101$ | VX | I | 323 | v3.0 |  |  | vrldnm | Vector Rotate Left Dword then AND with Mask |
| $000100 \ldots \ldots . . . . . . . . . .0000000110$ | VC | I | 306 | v2.03 |  |  | vcmpequb[.] | Vector Compare Equal Unsigned Byte |
|  | VC | I | 306 | v2.03 |  |  | vcmpequh[.] | Vector Compare Equal Unsigned Hword |
| $000100 \ldots \ldots . \ldots . . . . . . .0010000110$ | VC | I | 307 | v2.03 |  |  | vcmpequw[.] | Vector Compare Equal Unsigned Word |
|  | VC | I | 332 | v2.03 |  |  | vcmpeqfp[.] | Vector Compare Equal To Floating-Point |
|  | VC | 1 | 332 | v2.03 |  |  | vcmpgefp[.] | Vector Compare Greater Than or Equal To Floating-Point |
| 000100 $\ldots \ldots . . . . . . . . .1000000110$ | VC | 1 | 310 | v2.03 |  |  | vcmpgtub[.] | Vector Compare Greater Than Unsigned Byte |
| $000100 \ldots \ldots . . . . . . . .1001000110$ | VC | 1 | 311 | v2.03 |  |  | vcmpgtuh[.] | Vector Compare Greater Than Unsigned Hword |
| 000100 $\ldots \ldots . \ldots . . . . . . . . . .1010000110$ | VC | I | 311 | v2.03 |  |  | vcmpgtuw[.] | Vector Compare Greater Than Unsigned Word |
| 000100 $\ldots \ldots . \ldots . . . . . . .1011000110$ | VC | I | 333 | v2.03 |  |  | vcmpgtfp[.] | Vector Compare Greater Than Floating-Point |
| 000100 ......... ..... 1100000110 | VC | 1 | 308 | v2.03 |  |  | vcmpgtsb[.] | Vector Compare Greater Than Signed Byte |
| 000100 $\ldots \ldots . \ldots . . . . . . . .1101000110$ | VC | 1 | 309 | v2.03 |  |  | vcmpgtsh[.] | Vector Compare Greater Than Signed Hword |
| 000100 $\ldots \ldots . \ldots . . . . . . . . . .1110000110$ | VC | 1 | 309 | v2.03 |  |  | vcmpgtsw[.] | Vector Compare Greater Than Signed Word |
| $000100 \ldots \ldots . . . . . . . . .1111000110$ | VC | I | 331 | v2.03 |  |  | vcmpbfp[.] | Vector Compare Bounds Floating-Point |
| $000100 \ldots \ldots . \ldots . \ldots .0000000111$ | VC | I | 312 | v3.0 |  |  | vcmpneb[.] | Vector Compare Not Equal Byte |
|  | VC | 1 | 313 | v3.0 |  |  | vcmpneh[.] | Vector Compare Not Equal Hword |
|  | VC | 1 | 314 | v3.0 |  |  | vcmpnew[.] | Vector Compare Not Equal Word |
| $000100 \ldots \ldots . \ldots . . . . . .0011000111$ | VC | 1 | 307 | v2.07 |  |  | vcmpequd[.] | Vector Compare Equal Unsigned Dword |
| $000100 \ldots \ldots . \ldots . . . . . .0100000111$ | VC | I | 312 | v3.0 |  |  | vcmpnezb[.] | Vector Compare Not Equal or Zero Byte |
| 000100 $\ldots \ldots . \ldots . . . . . . . .0101000111$ | VC | 1 | 313 | v3.0 |  |  | vcmpnezh[.] | Vector Compare Not Equal or Zero Hword |
|  | VC | I | 314 | v3.0 |  |  | vcmpnezw[.] | Vector Compare Not Equal or Zero Word |

Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 3 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 4 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 5 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 6 of 17)

| Instruction $^{\mathbf{1}}$111111 11112 $22222 \quad 222233$  <br> $012345 \quad 67890$ 12345 67890 $12345 \quad 678901$ |  | $\begin{aligned} & \text { Y } \\ & \text { on } \\ & \text { on } \end{aligned}$ | $\begin{aligned} & 0 \\ & 0 \\ & 00 \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \frac{0}{\omega} \\ & \frac{N}{0} \\ & \hline \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $011110 \ldots . . . . . . . . . . . . . . . . . . .1001$. | MDS | 1 | 103 | PPC |  | SR | rldcr[.] | Rotate Left Dword then Clear Right |
| $011111 \ldots 1 . \ldots . . .1 . . .10000000001$ | X | I | 85 | P1 |  |  | cmp | Compare |
| 011111...1........ 00001000001 | X | I | 86 | P1 |  |  | cmpl | Compare Logical |
| 011111.......11 1111100100 000001 | X | I | 121 | v3.0 |  |  | setb | Set Boolean |
| $011111 \ldots 1 . \ldots . . . . . .10011000001$ | X | I | 87 | v3.0 |  |  | cmprb | Compare Ranged Byte |
| $011111 \ldots 11 \ldots . . . . . .100111000001$ | X | I | 88 | v3.0 |  |  | cmpeqb | Compare Equal Byte |
| $011111 \ldots 1111111111111001000000$ | X | I | 119 | v3.0 |  |  | mcrxrx | Move XER to CR Extended |
| 011111............ 00000001001 | X | I | 89 | P1 |  |  | tw | Trap Word |
| 011111............. 00010001001 | X | 1 | 90 | PPC |  |  | td | Trap Dword |
| 011111............ 00000001101 | X | I | 249 | v2.03 |  |  | Ivsl | Load Vector for Shift Left |
| 011111............ 00001001101 | X | 1 | 249 | v2.03 |  |  | Ivsr | Load Vector for Shift Right |
| 011111......... .... 10010001101 | X | 11 | 864 | v3.0 |  |  | Iwat | Load Word ATomic |
| 011111.............. 10011001101 | X | 11 | 864 | v3.0 |  |  | Idat | Load Dword ATomic |
| 011111......... ..... 10110001101 | X | 11 | 866 | v3.0 |  |  | stwat | Store Word ATomic |
| $011111 \ldots \ldots . . . . . . . . .10111001101$ | X | 11 | 866 | v3.0 |  |  | stdat | Store Dword ATomic |
| $0111111111 . \ldots . . . . . .11000001101$ | X | 11 | 858 | v3.0 |  |  | copy | Copy |
| 01111111111111111111111010001101 | X | 11 | 860 | v3.0 |  |  | cp_abort | CP_Abort |
| 011111 1111. ..... $\ldots$.... 1110000110. | X | 11 | 859 | v3.0 |  |  | paste[.] | Paste |
| 011111............. 00000001111 | X | 1 | 244 | v2.03 |  |  | Ivebx | Load Vector Element Byte Indexed |
|  | X | 1 | 244 | v2.03 |  |  | Ivehx | Load Vector Element Hword Indexed |
|  | X | I | 245 | v2.03 |  |  | Ivewx | Load Vector Element Word Indexed |
| 011111............. 00011001111 | X | 1 | 245 | v2.03 |  |  | Ivx | Load Vector Indexed |
|  | X | I | 247 | v2.03 |  |  | stvebx | Store Vector Element Byte Indexed |
| $011111 \ldots . . . . . . . . . . . .00101001111$ | X | I | 247 | v2.03 |  |  | stvehx | Store Vector Element Hword Indexed |
|  | X | I | 248 | v2.03 |  |  | stvewx | Store Vector Element Word Indexed |
| $011111 \ldots \ldots . . . . . . . . .00111001111$ | X | 1 | 248 | v2.03 |  |  | stvx | Store Vector Indexed |
|  | X | \| | 245 | v2.03 |  |  | Ivx\| | Load Vector Indexed Last |
| $011111 \ldots . . . . . . . . . .01111001111$ | X | I | 248 | v2.03 |  |  | stvx\| | Store Vector Indexed Last |
| 011111.......... $\ldots . . .0 .000001000$. | XO | 1 | 71 | P1 |  | SR | subfc[0][.] | Subtract From Carrying |
|  | XO | I | 70 | PPC |  | SR | subf[0][.] | Subtract From |
| 011111......... 11111.001101000. | XO | I | 73 | P1 |  | SR | neg[0][.] | Negate |
| 011111.............. . 010001000. | XO | 1 | 72 | P1 |  | SR | subfe[0][.] | Subtract From Extended |
| 011111......... 11111.011001000. | XO | I | 73 | P1 |  | SR | subfze[0][.] | Subtract From Zero Extended |
| 011111......... 11111.011101000. | XO | 1 | 72 | P1 |  | SR | subfme[0][.] | Subtract From Minus One Extended |
| 011111......... $\ldots$.... 1000001001. | XO | I | 80 | PPC |  | SR | mulhdu[.] | Multiply High Dword Unsigned |
| $\begin{array}{llllll}011111 & \ldots . . & \ldots . . & \ldots . . & 10010 & 01001 .\end{array}$ | XO | I | 80 | PPC |  | SR | mulhd[.] | Multiply High Dword |
|  | XO | 1 | 80 | PPC |  | SR | mulld[0][.] | Multiply Low Dword |
| 011111 $\ldots . . . . . . . . . . .110001001$. | XO | I | 83 | v2.06 |  | SR | divdeu[0][.] | Divide Dword Extended Unsigned |
| 011111............... 110101001. | XO | I | 83 | V2.06 |  | SR | divde[0][.] | Divide Dword Extended |
|  | XO | I | 82 | PPC |  | SR | divdu[0][.] | Divide Dword Unsigned |
| 011111 $\ldots . . . . . . . . . . .111101001$. | XO | I | 82 | PPC |  | SR | divd[0][.] | Divide Dword |
| 011111............ 01000010011 | X | I | 84 | v3.0 |  |  | modud | Modulo Unsigned Dword |
| $011111 \ldots \ldots . . . . . . . .11000010011$ | X | 1 | 84 | v3.0 |  |  | modsd | Modulo Signed Dword |
|  | XO | 1 | 71 | P1 |  | SR | addc[0][.] | Add Carrying |
| 011111......... $\ldots$.... 10010010101 | XO | I | 110 | v2.06 |  |  | addg6s | Add \& Generate Sixes |
|  | XO | I | 72 | P1 |  | SR | adde[0][.] | Add Extended |
|  | XO | 1 | 73 | P1 |  | SR | addze[0][.] | Add to Zero Extended |
|  | XO | 1 | 72 | P1 |  | SR | addme[0][.] | Add to Minus One Extended |
| 011111.............. 100001010. | XO | I | 70 | P1 |  | SR | add[0][.] | Add |
| $011111 \ldots \ldots . . . . . . . .1000001011$. | XO | I | 74 | PPC |  | SR | mulhwu[.] | Multiply High Word Unsigned |
| 011111......... $\ldots . . .1001001011$. | XO | 1 | 74 | PPC |  | SR | mulhw[.] | Multiply High Word |
| 011111 $\ldots .$. $\ldots .$. $\ldots .$. 0111 01011. | XO | 1 | 74 | P1 |  | SR | mullw[0][.] | Multiply Low Word |
| 0111111.............. . 110001011. | XO | I | 77 | v2.06 |  | SR | divweu[0][.] | Divide Word Extended Unsigned |
|  | XO | 1 | 77 | v2.06 |  | SR | divwe[0][.] | Divide Word Extended |
|  | XO | + | 75 | PPC |  | SR | divwu[0][.] | Divide Word Unsigned |
| $011111 \ldots . . . . . . . . . . . .111101011$. | XO | 1 | 75 | PPC |  | SR | divw[0][.] | Divide Word |

Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 7 of 17)

| Instruction  <br> 1  <br> 1 11111 <br> 11112 22222222233 <br> 01234567890 12345 <br> 67890 12345678901 |  | $\begin{aligned} & \text { Yo } \\ & \text { O } \end{aligned}$ | $\begin{aligned} & \stackrel{0}{\widetilde{\circ}} \\ & \hline \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \text {. } \frac{0}{n} \\ & \stackrel{\omega}{\omega} \end{aligned}$ |  | $\begin{array}{\|c} r_{0}^{0} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \dot{0} \\ \hline \end{array}$ |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 011111 ..... ..... ..... 01000010111 | X | 1 | 76 | v3.0 |  |  | moduw | Modulo Unsigned Word |
| $011111 . . . .1 . . . . . . . . .11000010111$ | X | 1 | 76 | v3.0 |  |  | modsw | Modulo Signed Word |
|  | XX1 | 1 | 485 | v2.07 |  |  | Ixsiwzx | Load VSX Scalar as Integer Word \& Zero Indexed |
| $011111 . \ldots . . . . . . . . . .00001001100$. | XX1 | 1 | 484 | v2.07 |  |  | Ixsiwax | Load VSX Scalar as Integer Word Algebraic Indexed |
| $\begin{array}{llllllll}0111111 & \ldots . . & \ldots . . & & 00100 & 01100\end{array}$ | XX1 | 1 | 501 | v2.07 |  |  | stxsiwx | Store VSX Scalar as Integer Word Indexed |
| $011111 . . . . . . . . . . . . . .0100001100$. | XX1 | 1 | 493 | v3.0 |  |  | Ixvx | Load VSX Vector Indexed |
|  | XX1 | 1 | 495 | v2.06 |  |  | Ixvdsx | Load VSX Vector Dword \& Splat Indexed |
|  | XX1 | 1 | 498 | v3.0 |  |  | lxwsx | Load VSX Vector Word \& Splat Indexed |
|  | XX1 | 1 | 511 | v3.0 |  |  | stxvx | Store VSX Vector Indexed |
| $011111 \ldots \ldots . . . . . . . .1000001100$. | XX1 | 1 | 486 | v2.07 |  |  | Ixsspx | Load VSX Scalar SP Indexed |
| 011111 $\ldots .$. $\ldots .$. $\ldots .$. 10010 01100. | XX1 | 1 | 481 | v2.06 |  |  | Ixsdx | Load VSX Scalar Dword Indexed |
|  | XX1 | 1 | 503 | v2.07 |  |  | stxsspx | Store VSX Scalar SP Indexed |
|  | XX1 | 1 | 499 | v2.06 |  |  | stxsdx | Store VSX Scalar Dword Indexed |
|  | XX1 | 1 | 497 | v2.06 |  |  | 1xw ${ }^{\text {dx }}$ | Load VSX Vector Word*4 Indexed |
| $011111 . . . .1 . . . . . . . . .1100101100$. | XX1 | 1 | 496 | v3.0 |  |  | Ixvh8x | Load VSX Vector Hword*8 Indexed |
|  | XX1 | 1 | 489 | v2.06 |  |  | \|xvd2x | Load VSX Vector Dword*2 Indexed |
| 011111 ..... ......... 1101101100. | XX1 | 1 | 488 | v3.0 |  |  | lxvb16x | Load VSX Vector Byte*16 Indexed |
| $\begin{array}{lllllllll}011111 & \ldots . . & \ldots . . & \ldots 1100 & 01100\end{array}$ | XX1 | 1 | 507 | v2.06 |  |  | stxvw4x | Store VSX Vector Word*4 Indexed |
|  | XX1 | 1 | 506 | v3.0 |  |  | stxvh8x | Store VSX Vector Hword*8 Indexed |
|  | XX1 | 1 | 505 | v2.06 |  |  | stxvd2x | Store VSX Vector Dword*2 Indexed |
| $\begin{array}{lllllllll}011111 & \ldots . . & \ldots . . & \ldots 111101100 .\end{array}$ | XX1 | 1 | 504 | v3.0 |  |  | stxvb16x | Store VSX Vector Byte*16 Indexed |
|  | XX1 | 1 | 490 | v3.0 |  |  | Ixv\| | Load VSX Vector with Length |
|  | XX1 | 1 | 492 | v3.0 |  |  | \|xv|l | Load VSX Vector Left-justified with Length |
|  | XX1 | I | 508 | v3.0 |  |  | stxvl | Store VSX Vector with Length |
| $\begin{array}{lllllllll}011111 & \ldots . . & \ldots . . & \ldots 110101101 .\end{array}$ | XX1 | 1 | 510 | v3.0 |  |  | stxvll | Store VSX Vector Left-justified with Length |
| $011111 . \ldots . . . . . . . . . . .1100001101$. | XX1 | 1 | 483 | v3.0 |  |  | 1xsibzx | Load VSX Scalar as Integer Byte \& Zero Indexed |
| $011111 . . . .1 . . . . . . . . .1100101101$. | XX1 | 1 | 483 | v3.0 |  |  | 1xsihzx | Load VSX Scalar as Integer Hword \& Zero Indexed |
|  | XX1 | 1 | 500 | v3.0 |  |  | stxsibx | Store VSX Scalar as Integer Byte Indexed |
|  | XX1 | I | 500 | v3.0 |  |  | stxsihx | Store VSX Scalar as Integer Hword Indexed |
| 011111 \|1111 |1111 .... 00100011101 | X | III | 1125 | v2.07 | P |  | msgsndp | Message Send Privileged |
| 011111 \||1|| |1111 .... 00101011101 | X | III | 1126 | v2.07 | P |  | msgclip | Message Clear Privileged |
| $0111111111111111 \ldots 00110011101$ | X | 111 | 1123 | v2.07 | H |  | msgsnd | Message Send |
| $0111111111111111 \ldots . .00111011101$ | X | 111 | 1124 | v2.07 | H |  | msgclr | Message Clear |
| 011111 ..... ..... ..... 01001011101 | X | 1 | 44 | v2.07 |  |  | mfbhrbe | Move From BHRB |
| 011111 \|1|1| ||11| |1||| 01101011101 | X | 1 | 44 | v2.07 |  |  | clirbhrb | Clear BHRB |
| $011111.1111\|111\| 1111 \mid 10101011101$ | X | 11 | 894 | v2.07 |  |  | tend. | Transaction End \& record |
| 011111 ...11 \||111||1||110110 011101 | X | 11 | 898 | v2.07 |  |  | tcheck | Transaction Check \& record |
| 011111 \|111. 11111 |111110111 011101 | X | 11 | 898 | v2.07 |  |  | tsr. | Transaction Suspend or Resume \& record |
| 011111.111. 11111 \|111110100 011101 | X | 11 | 893 | v2.07 |  |  | tbegin. | Transaction Begin \& record |
|  | X | 11 | 896 | v2.07 |  |  | tabortwc. | Transaction Abort Word Conditional \& record |
|  | X | 11 | 897 | v2.07 |  |  | tabortdc. | Transaction Abort Dword Conditional \& record |
| 011111 ..... .......... 11010011101 | X | 11 | 896 | v2.07 |  |  | tabortwci. | Transaction Abort Word Conditional Immediate \& record |
| 011111 ..... ..... ..... 11011011101 | X | 11 | 897 | v2.07 |  |  | tabortdci. | Transaction Abort Dword Conditional Immediate \& record |
| $011111\|1\| 11 . \ldots . .\|1\| 1 \mid 11100011101$ | X | 11 | 895 | v2.07 |  |  | tabort. | Transaction Abort \& record |
| 011111 \|1111 ..... |111| 11101 011101 | X | 111 | 969 | v2.07 | P |  | treclaim. | Transaction Reclaim \& record |
| 011111 \||||| ||||| ||||| 11111011101 | X | 111 | 970 | v2.07 | P |  | trechkpt. | Transaction Recheckpoint \& record |
|  | A | 1 | 90 | v2.03 |  |  | isel | Integer Select |
| $011111 \ldots . .0 \ldots \ldots .100100100001$ | XFX | 1 | 120 | P1 |  |  | mtcrf | Move To CR Fields |
| 011111 ..... 1....... 00100100001 | XFX | I | 120 | v2.01 |  |  | mtocif | Move To One CR Field |
| $011111 \ldots . .1111 .1111100100100101$ | X | 111 | 977 | P1 | P |  | mtmsr | Move To MSR |
| 011111 .... \||11. |1111 00101100101 | X | 111 | 978 | PPC | P |  | mtmsrd | Move To MSR Dword |
| $0111111111111111 \ldots . .01000100101$ | X | III | 1038 | v2.03 | P | 64 | tlbiel | TLB Invalidate Entry Local |
| $0111111111.11111 \ldots 010001100101$ | X | 111 | 1034 | P1 | H | 64 | tlbie | TLB Invalidate Entry |
| 011111 \||1|| ||1|| |1|1| 01010100101 | X | III | 1031 | v3.0 | P |  | slbsync | SLB Synchronize |
| $011111 . . . .11111$..... 01100100101 | X | 111 | 1029 | v2.00 | P |  | slbmte | SLB Move To Entry |
| 011111 \|111| |1111..... 01101100101 | X | III | 1024 | PPC | P |  | slbie | SLB Invalidate Entry |

Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 8 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 9 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 10 of 17)

| Instruction $^{\mathbf{1}}$11 11111 $11112 \quad 22222222233$ <br> $012345 \quad 67890$ 1234567890 $12345 \quad 678901$ |  | $\begin{aligned} & \text { y } \\ & \text { on } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & \underset{\sim}{0} \end{aligned}$ |  |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 011111 $\ldots \ldots . . . . .1111 \mid 1110011010$. | X | I | 94 | P1 |  | SR | extsh[.] | Extend Sign Hword |
| 011111..... $\ldots . . .1\|11\| 1110111010$. | X | 1 | 94 | PPC |  | SR | extsb[.] | Extend Sign Byte |
|  | X | 1 | 98 | PPC |  | SR | extsw[.] | Extend Sign Word |
|  | X | I | 108 | PPC |  | SR | sld[.] | Shift Left Dword |
| 011111 $\ldots . .$. $\ldots . .$. 1000011011. | X | I | 108 | PPC |  | SR | srd[.] | Shift Right Dword |
|  | X | I | 93 | P1 |  | SR | and[.] | AND |
| $011111 \ldots \ldots . . . . . . . . .0000111100$. | X | I | 94 | P1 |  | SR | andc[.] | AND with Complement |
| 011111 $\ldots . .$. $\ldots . .$. 0001111100 | X | I | 94 | P1 |  | SR | nor[.] | NOR |
| 011111.............. 00111111001 | X | 1 | 99 | v2.06 |  |  | bpermd | Bit Permute Dword |
|  | X | 1 | 94 | P1 |  | SR | eqv[.] | Equivalent |
| 011111 $\ldots .$. $\ldots . .$. $\ldots 100111100$ | X | 1 | 93 | P1 |  | SR | xor[.] | XOR |
| 011111 $\ldots .$. $\ldots . .$. $\ldots 110011100$ | X | I | 94 | P1 |  | SR | orc[.] | OR with Complement |
| 011111 $\ldots .$. $\ldots .$. $\ldots 110111100$ | X | I | 93 | P1 |  | SR | or[.] | OR |
| $011111 . . . . . . . . . . . . . ~ 0111011100$. | X | I | 93 | P1 |  | SR | nand[.] | NAND |
|  | X | I | 96 | v2.05 |  |  | cmpb | Compare Bytes |
| $011111111 . .111111111100000111101$ | X | 11 | 880 | v3.0 |  |  | wait | Wait for Interrupt |
|  | D | I | 51 | P1 |  |  | IWz | Load Word \& Zero |
|  | D | 1 | 51 | P1 |  |  | IWzu | Load Word \& Zero with Update |
|  | D | I | 48 | P1 |  |  | lbz | Load Byte \& Zero |
|  | D | I | 48 | P1 |  |  | lbzu | Load Byte \& Zero with Update |
|  | D | I | 57 | P1 |  |  | stw | Store Word |
| 100101 .......................... | D | I | 57 | P1 |  |  | stwu | Store Word with Update |
|  | D | I | 55 | P1 |  |  | stb | Store Byte |
|  | D | I | 55 | P1 |  |  | stbu | Store Byte with Update |
|  | D | I | 49 | P1 |  |  | lhz | Load Hword \& Zero |
| 101001 .......... .................. | D | 1 | 49 | P1 |  |  | Ihzu | Load Hword \& Zero with Update |
|  | D | 1 | 50 | P1 |  |  | Iha | Load Hword Algebraic |
|  | D | । | 50 | P1 |  |  | Ihau | Load Hword Algebraic with Update |
|  | D | I | 56 | P1 |  |  | sth | Store Hword |
|  | D | I | 56 | P1 |  |  | sthu | Store Hword with Update |
|  | D | I | 63 | P1 |  |  | Imw | Load Multiple Word |
|  | D | I | 63 | P1 |  |  | stmw | Store Multiple Word |
|  | D | I | 142 | P1 |  |  | Ifs | Load Floating Single |
|  | D | I | 142 | P1 |  |  | Ifsu | Load Floating Single with Update |
|  | D | I | 143 | P1 |  |  | Ifd | Load Floating Double |
|  | D | I | 143 | P1 |  |  | Ifdu | Load Floating Double with Update |
|  | D | 1 | 146 | P1 |  |  | stfs | Store Floating Single |
|  | D | I | 146 | P1 |  |  | stfsu | Store Floating Single with Update |
|  | D | I | 147 | P1 |  |  | stfd | Store Floating Double |
| 110111 ......................... | D | I | 147 | P1 |  |  | stfdu | Store Floating Double with Update |
|  | DQ | I | 59 | v2.03 |  |  | Iq | Load Qword |
|  | DS | I | 150 | v2.05 |  |  | Ifdp | Load Floating Double Pair |
|  | DS | 1 | 481 | v3.0 |  |  | Ixsd | Load VSX Scalar Dword |
|  | DS | I | 486 | v3.0 |  |  | Ixssp | Load VSX Scalar Single |
|  | DS | I | 53 | PPC |  |  | Id | Load Dword |
|  | DS | I | 53 | PPC |  |  | Idu | Load Dword with Update |
|  | DS | I | 52 | PPC |  |  | Iwa | Load Word Algebraic |
| 111011.............. 001000010. | Z22 | I | 222 | v2.05 |  |  | dscli[.] | DFP Shift Significand Left Immediate |
| 111011.............. . 001100010. | Z22 | I | 222 | v2.05 |  |  | dscri[.] | DFP Shift Significand Right Immediate |
|  | Z22 | I | 202 | v2.05 |  |  | dtstdc | DFP Test Data Class |
| 111011 $\ldots .11 \ldots . . . . . . .0111000101$ | Z22 | 1 | 202 | v2.05 |  |  | dtstdg | DFP Test Data Group |
| 111011 $\ldots \ldots . . . . . . . . . . . ~ 0000000010 . ~$ | X | I | 195 | v2.05 |  |  | dadd[.] | DFP Add |
| 111011............. 0000100010. | X | I | 197 | v2.05 |  |  | dmul[.] | DFP Multiply |
| 111011...11 $\ldots . . . \ldots . .00100000101$ | X | I | 201 | v2.05 |  |  | dcmpo | DFP Compare Ordered |
| 111011 $\ldots .11 \ldots \ldots . . . . .100101000101$ | X | I | 203 | v2.05 |  |  | dtstex | DFP Test Exponent |
| 111011..... 11111..... 0100000010. | X | I | 215 | v2.05 |  |  | dctdp[.] | DFP Convert To DFP Long |

Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 11 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 12 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 13 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 14 of 17)

| Instruction ${ }^{1}$ |  |  |  |  |  |  |  | $\begin{aligned} & \mathbb{\otimes} \\ & \stackrel{\sim}{0} \end{aligned}$ |  |  |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 012345 | $678901$ | $\begin{aligned} & 111111 \\ & 12345 \end{aligned}$ |  | $\begin{aligned} & 22222 \\ & 12345 \end{aligned}$ |  |  |  |  |  |  |  |  |  |  |
| 111100 | ..... 1 | 11111 | ..... | 11000 | 1001. | XX2 | I | 676 |  | v2.06 |  |  | xvcvdpsp | VSX Vector Convert DP to SP |
| 111100 | .... 1 | 11111 | ..... | 11001 | 1001 | XX2 | I | 660 |  | v2.06 |  |  | xvabssp | VSX Vector Absolute SP |
| 111100 | $\ldots$ | 11111 | ..... | 11010 | 1001. | XX2 | I | 729 |  | v2.06 |  |  | xvnabssp | VSX Vector Negative Absolute SP |
| 111100 | $\ldots 1$ | 11111 | ..... | 11011 | 1001. | XX2 | I | 730 |  | v2.06 |  |  | xvnegsp | VSX Vector Negate SP |
| 111100 | $\ldots 1$ | 11111 | ..... | 11100 | 1001. | XX2 | I | 686 |  | v2.06 |  |  | xvcvspdp | VSX Vector Convert SP to DP |
| 111100 | .... 1 | 11111 | ..... | 11101 | 1001. | XX2 | I | 660 |  | v2.06 |  |  | xvabsdp | VSX Vector Absolute DP |
| 111100 | $\ldots$ | 11111 | ..... | 11110 | 1001. | XX2 | I | 729 |  | v2.06 |  |  | xvnabsdp | VSX Vector Negative Absolute DP |
| 111100 | .... 1 | 11111 | ..... | 11111 | 1001. | XX2 | I | 730 |  | v2.06 |  |  | xvnegdp | VSX Vector Negate DP |
| 111100 | . .11 | . .... | ..... | 00111 | 101.1 | XX3 | I | 65 |  | v2.06 |  |  | xstdivdp | VSX Scalar Test for software Divide DP |
| 111100 | . 111 | ..... | ..... | 01011 | 101..1 | XX3 | I | 76 |  | v2.06 |  |  | xutdivsp | VSX Vector Test for software Divide SP |
| 111100 | $\ldots 11$. | ..... | ..... | 01111 | 101. 1 | XX3 | I | 76 |  | v2.06 |  |  | xutdivdp | VSX Vector Test for software Divide DP |
| 111100 | .... | ..... | ..... | 1101. | 101. | XX2 | I | 76 |  | v3.0 |  |  | xvtstdcsp | VSX Vector Test Data Class SP |
| 111100 | ...... | ..... | ..... | 1111. | 101. | XX2 | I | 76 |  | v3.0 |  |  | xvtstdcdp | VSX Vector Test Data Class DP |
| 111100 | ..... 1 | 11111 | ..... | 00000 | 1010 | XX2 | I | 64 |  | v2.07 |  |  | xsisqritesp | VSX Scalar Reciprocal Square Root Estimate SP |
| 111100 | .... 1 | 11111 | ..... | 00001 | 1010 | XX2 | I | 63 |  | v2.07 |  |  | xsresp | VSX Scalar Reciprocal Estimate SP |
| 111100 | .... 1 | 11111 | ..... | 00100 | 1010. | XX2 | I | 64 |  | v2.06 |  |  | xsrsqritedp | VSX Scalar Reciprocal Square Root Estimate DP |
| 111100 | $\ldots$ | 11111 | ..... | 00101 | 1010 | XX2 | I | 63 |  | v2.06 |  |  | xsredp | VSX Scalar Reciprocal Estimate DP |
| 111100 | $\ldots 111$ | 11111 | ..... | 00110 | 1010.1 | XX2 | I | 65 |  | v2.06 |  |  | xstsqrtdp | VSX Scalar Test for software Square Root DP |
| 111100 | $\ldots 1$ | 11111 | ..... | 01000 | 1010 | XX2 | I | 75 |  | v2.06 |  |  | xvrsqritesp | VSX Vector Reciprocal Square Root Estimate SP |
| 111100 | $\ldots$ | 11111 | ..... | 01001 | 1010 | XX2 | I | 74 |  | v2.06 |  |  | xvresp | VSX Vector Reciprocal Estimate SP |
| 111100 | $\ldots 111$ | 11111 | ..... | 01010 | 1010.1 | XX2 | I | 76 |  | v2.06 |  |  | xvtsqrtsp | VSX Vector Test for software Square Root SP |
| 111100 | .... 1 | 11111 | ..... | 01100 | 1010. | XX2 | I | 75 |  | v2.06 |  |  | xvrsqritedp | VSX Vector Reciprocal Square Root Estimate DP |
| 111100 | $\ldots$ | 11111 | ..... | 01101 | 1010 | XX2 | 1 | 74 |  | v2.06 |  |  | xvredp | VSX Vector Reciprocal Estimate DP |
| 111100 | ...11 1 | 11111 | ..... | 01110 | 1010.1 | XX2 | I | 76 |  | v2.06 |  |  | xvtsqrtdp | VSX Vector Test for software Square Root DP |
| 111100 | .... | ..... | ..... | 10010 | 1010.1 | XX2 | I | 65 |  | v3.0 |  |  | xststdcsp | VSX Scalar Test Data Class SP |
| 111100 | ..... | ..... | ..... | 10110 | 1010.1 | XX2 | I | 65 |  | v3.0 |  |  | xststdcdp | VSX Scalar Test Data Class DP |
| 111100 | ..... 1 | 11111 | ..... | 00000 | 1011. | XX2 | I | 64 |  | v2.07 |  |  | xssqrisp | VSX Scalar Square Root SP |
| 111100 | $\ldots$ | 11111 | ..... | 00100 | 1011 | XX2 | 1 | 64 |  | v2.06 |  |  | xssqrtdp | VSX Scalar Square Root DP |
| 111100 | ..... 1 | 11111 | ..... | 00110 | 1011 | XX2 | I | 63 |  | v2.06 |  |  | xsrdpic | VSX Scalar Round DP to Integral using Current rounding mode |
| 111100 | $\ldots$ | 11111 | ..... | 01000 | 1011. | XX2 | I | 75 |  | v2.06 |  |  | xvsqrisp | VSX Vector Square Root SP |
| 111100 | $\ldots$ | 11111 | ..... | 01010 | 1011 | XX2 | I | 75 |  | v2.06 |  |  | xvrspic | VSX Vector Round SP to Integral using Current rounding mode |
| 111100 | ..... 1 | 11111 | ..... | 01100 | 1011. | XX2 | I | 75 |  | v2.06 |  |  | xvsqrtdp | VSX Vector Square Root DP |
| 111100 | $\ldots$ | 11111 | ..... | 01110 | 1011. | XX2 | I | 74 |  | v2.06 |  |  | xvrdpic | VSX Vector Round DP to Integral using Current rounding mode |
| 111100 | $\ldots$ | 11111 | ..... | 10000 | 1011. | XX2 | I | 53 |  | v2.07 |  |  | xscvdpspn | VSX Scalar Convert DP to SP Non-signalling |
| 111100 | ..... 1 | 11111 | ..... | 10100 | 1011. | XX2 | I | 56 |  | v2.07 |  |  | xscvspdpn | VSX Scalar Convert SP to DP Non-signalling |
| 111100 | .... 0 | 00000 | ..... | 10101 | 1011.1 | XX2 | I | 65 |  | v3.0 |  |  | xsxexpdp | VSX Scalar Extract Exponent DP |
| 111100 | $\ldots$ | 00001 | ..... | 10101 | 1011.1 | XX2 | I | 65 |  | v3.0 |  |  | xsxsigdp | VSX Scalar Extract Significand DP |
| 111100 | ..... 1 | 10000 | ..... | 10101 | 1011. | XX2 | I | 54 |  | v3.0 |  |  | xscvhpdp | VSX Scalar Convert HP to DP |
| 111100 | $\ldots$ | 10001 | ..... | 10101 | 1011. | XX2 | I | 53 |  | v3.0 |  |  | xscvdphp | VSX Scalar Convert DP to HP |
| 111100 | $\ldots$ | 00000 | ..... | 11101 | 1011. | XX2 | I | 76 |  | v3.0 |  |  | xvxexpdp | VSX Vector Extract Exponent DP |
| 111100 | $\ldots$ | 00001 | ..... | 11101 | 1011. | XX2 | I | 76 |  | v3.0 |  |  | xvxsigdp | VSX Vector Extract Significand DP |
| 111100 | .... 0 | 00111 | ..... | 11101 | 1011. | XX2 | I | 76 |  | v3.0 |  |  | xxbrh | VSX Vector Byte-Reverse Hword |
| 111100 | .... 0 | 01000 | ..... | 11101 | 1011. | XX2 | I | 76 |  | v3.0 |  |  | xvxexpsp | VSX Vector Extract Exponent SP |
| 111100 | .... 0 | 01001 | ..... | 11101 | 1011. | XX2 | I | 76 |  | v3.0 |  |  | xvxsigsp | VSX Vector Extract Significand SP |
| 111100 | ..... 0 | 01111 | ..... | 11101 | 1011. | XX2 | I | 76 |  | v3.0 |  |  | xxbrw | VSX Vector Byte-Reverse Word |
| 111100 | ..... 1 | 10111 | ..... | 11101 | 1011. | XX2 | I | 76 |  | v3.0 |  |  | xxbrd | VSX Vector Byte-Reverse Dword |
| 111100 | ..... | 11000 | ..... | 11101 | 1011. | XX2 | I | 68 |  | v3.0 |  |  | xvcvhpsp | VSX Vector Convert HP to SP |
| 111100 | ..... | 11001 | ..... | 11101 | 1011 | XX2 | I | 68 |  | v3.0 |  |  | xvcvsphp | VSX Vector Convert SP to HP |
| 111100 | .... | 11111 | ..... | 11101 | 1011. | XX2 | I | 76 |  | v3.0 |  |  | xxbrq | VSX Vector Byte-Reverse Qword |
| 111100 | .... | ..... | ...... | 11100 | 10110 | XX1 | I | 57 |  | v3.0 |  |  | xsiexpdp | VSX Scalar Insert Exponent DP |
| 111100 | .... | ..... | ..... | ..... | 11. | XX4 | I | 77 |  | v2.06 |  |  | xxsel | VSX Vector Select |
| 111101 | ..... | ..... | ..... | ..... | $\ldots$ | DS | I | 15 |  | v2.05 |  |  | sttdp | Store Floating Double Pair |
| 111101 | .... | ..... | ..... | ..... | $\ldots$ | DS | I | 49 |  | v3.0 |  |  | stxsd | Store VSX Scalar Dword |
| 111101 | ..... | ..... | ..... | ..... | $\ldots 11$ | DS | I | 50 |  | v3.0 |  |  | stxssp | Store VSX Scalar SP |
| 111101 | ..... | ..... | ..... | ..... | $\ldots 001$ | DQ | I | 49 |  | v3.0 |  |  | Ixv | Load VSX Vector |
| 111101 | ..... | ..... | ..... | ..... | . . 101 | DQ | I | 50 |  | v3.0 |  |  | stzv | Store VSX Vector |

Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 15 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 16 of 17)


Figure 87. Power ISA Instruction Set Sorted by Opcode (Sheet 17 of 17)

1. Key to Instruction column (primary and extended opcode bits shaded in gray).

1 Instruction bit that corresponds to a reserved field, must have a value of 0 , otherwise invalid form. Instruction bit that corresponds to an operand bit, may have a value of either 0 or 1.
$0 \quad$ Instruction bit having a value 0.
1 Instruction bit having a value 1.
2. Key to Version column.

| P1 | Instruction introduced in the POWER Architecture. |
| :---: | :--- |
| P2 | Instruction introduced in the POWER2 Architecture. |
| PPC | Instruction introduced in the PowerPC Architecture prior to v2.00. |
| v2.00 | Instruction introduced in the PowerPC Architecture Version 2.00. |
| v2.01 | Instruction introduced in the PowerPC Architecture Version 2.01. |
| v2.02 | Instruction introduced in the PowerPC Architecture Version 2.02. |
| v2.03 | Instruction introduced in the Power ISA Architecture Version 2.03. |
| v2.04 | Instruction introduced in the Power ISA Architecture Version 2.04. |
| v2.05 | Instruction introduced in the Power ISA Architecture Version 2.05. |
| v2.06 | Instruction introduced in the Power ISA Architecture Version 2.06. |
| v2.07 | Instruction introduced in the Power ISA Architecture Version 2.07. |
| v3.0 | Instruction introduced in the Power ISA Architecture Version 3.00. |

3. Key to Privilege column.

P Denotes an instruction that is treated as privileged.
O Denotes an instruction that is treated as privileged or nonprivileged (or hypervisor, for mtspr), depending on the SPR or PMR number.
PI Denotes an instruction that is illegal in privileged state.
H Denotes an instruction that can be executed only in hypervisor state
4. Key to Mode Dependency column.

Except as described below and in Section 1.11.3, "Effective Address Calculation", in Book I, all instructions are independent of whether the processor is in 32 -bit or 64 -bit mode.

CT If the instruction tests the Count Register, it tests the low-order 32 bits in 32-bit mode and all 64 bits in 64-bit mode.
SR The setting of status registers (such as XER and CRO) is mode-dependent.
32 The instruction can be executed only in 32-bit mode.
64 The instruction can be executed only in 64-bit mode.

## Appendix E. Power ISA Instruction Set Sorted by Version

This appendix lists all the instructions in the Power ISA, sorted in reverse order by ISA version.

| Instruction ${ }^{\mathbf{1}}$ 111111111222222222233 01234567890123456789012345678901 |  | $\begin{aligned} & \text { 등 } \\ & \text { © } \end{aligned}$ | $\begin{aligned} & \stackrel{0}{ஜ} \\ & \stackrel{0}{2} \end{aligned}$ | $\begin{aligned} & N \\ & \stackrel{0}{0} \\ & \stackrel{N}{0} \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | DX | 1 | 69 | v3.0 |  |  | addpcis | Add PC Immediate Shitted |
| $000100 \ldots . .00111 \ldots \ldots 1.110000001$ | VX | I | 352 | v3.0 |  |  | bcdcfn. | Decimal Convert From National \& record |
| $000100 \ldots . .00010 \ldots \ldots 1.110000001$ | VX | 1 | 356 | v3.0 |  |  | bcdofsq. | Decimal Convert From Signed Qword \& record |
| $000100 \ldots . .00110 \ldots \ldots 1.110000001$ | VX | I | 353 | v3.0 |  |  | bcdicz. | Decimal Convert From Zoned \& record |
| $000100 \ldots \ldots . . . . . . .01101000001$ | VX | I | 358 | v3.0 |  |  | bcdcpsgn. | Decimal CopySign \& record |
| $000100 \ldots \ldots 00101 \ldots \ldots 1 / 110000001$ | VX | 1 | 354 | v3.0 |  |  | bcdctn. | Decimal Convert To National \& record |
| $000100 \ldots . .00000 \ldots \ldots 1 / 110000001$ | VX | I | 356 | v3.0 |  |  | bcdctsq. | Decimal Convert To Signed Qword \& record |
| $000100 \ldots . .00100 \ldots \ldots 1.110000001$ | VX | 1 | 355 | v3.0 |  |  | bcdctz. | Decimal Convert To Zoned \& record |
| $000100 \ldots \ldots . . . . .1 . . .1 .0111000001$ | VX | 1 | 359 | v3.0 |  |  | bcds. | Decimal Shitt \& record |
| $000100 \ldots \ldots 11111 \ldots \ldots 1.110000001$ | VX | 1 | 358 | v3.0 |  |  | bcdsetsgn. | Decimal Set Sign \& record |
| $000100 \ldots \ldots . . . . . . . .1 .111000001$ | VX | 1 | 361 | v3.0 |  |  | bcdsr. | Decimal Shitt \& Round \& record |
|  | VX | 1 | 362 | v3.0 |  |  | bcdtrunc. | Decimal Truncate \& record |
| 000100 .............. 1/010 000001 | VX | 1 | 360 | v3.0 |  |  | bcdus. | Decimal Unsigned Shitt \& record |
| $000100 \ldots \ldots . . . . . . . . .1 / 101000001$ | VX | 1 | 363 | v3.0 |  |  | bcdutrunc. | Decimal Unsigned Truncate \& record |
| $011111 \ldots .11 \ldots \ldots . . .00111000001$ | X | I | 88 | v3.0 |  |  | cmpeqb | Compare Equal Byte |
| $011111 . .1 .1 . . . .1 . . .00110000001$ | X | I | 87 | v3.0 |  |  | cmprb | Compare Ranged Byte |
| $011111 \ldots \ldots . . . .11\|1\| 1000111010$. | X | 1 | 98 | v3.0 |  |  | cnttzd[.] | Count Trailing Zeros Dword |
| $011111 \ldots \ldots . . . .111111000011010$ | X | 1 | 95 | v3.0 |  |  | cnttzw[] | Count Trailing Zeros Word |
| $0111111111 . . . . .1 . . .11000001101$ | X | 11 | 858 | v3.0 |  |  | copy | Copy |
| $01111111111111111111 \mid 11010001101$ | X | 11 | 860 | v3.0 |  |  | cp_abort | CP_Abort |
| 011111 ..... \|11.. |1||| 10111100111 | X | 1 | 79 | v3.0 |  |  | darn | Deliver A Random Number |
| 111011...11 .......... 10101000111 | X | 1 | 204 | v3.0 |  |  | dtstsfi | DFP Test Significance Immediate |
| 111111 ...11 ......... $1010100011 /$ | X | 1 | 204 | v3.0 |  |  | dtstsfiq | DFP Test Significance Immediate Quad |
|  | XS | 1 | 109 | v3.0 |  |  | extswsli[.] | Extend Sign Word \& Shift Left Immediate |
| $011111 \ldots . . . . . . . . . .10011001101$ | X | 11 | 864 | v3.0 |  |  | Idat | Load Dword ATomic |
|  | X | 1 | 54 | v3.0 | Pl |  | Idmx | Load Dword Monitored Indexed |
| $011111 \ldots \ldots . . . .10 . .10010001101$ | X | 11 | 864 | v3.0 |  |  | Iwat | Load Word ATomic |
| 111001 ..... ..... ..... ..... .... 10 | DS | 1 | 481 | v3.0 |  |  | Ixsd | Load VSX Scalar Dword |
|  | XX1 | I | 483 | v3.0 |  |  | \|xsibzx | Load VSX Scalar as Integer Byte \& Zero Indexed |
|  | XX1 | 1 | 483 | v3.0 |  |  | \|xsihzx | Load VSX Scalar as Integer Hword \& Zero Indexed |
|  | DS | I | 486 | v3.0 |  |  | \|xssp | Load VSX Scalar Single |
|  | DQ | I | 493 | v3.0 |  |  | IxV | Load VSX Vector |
| 011111 ..... ..... .... 1101101100. | XX1 | I | 488 | v3.0 |  |  | \|xvb16x | Load VSX Vector Byte*16 Indexed |
| $011111 \ldots \ldots . . . . . . . .1100101100$. | XX1 | 1 | 496 | v3.0 |  |  | Ixvh8x | Load VSX Vector Hword*8 Indexed |
|  | XX1 | 1 | 490 | v3.0 |  |  | \|xv| | Load VSX Vector with Length |
| $011111 \ldots . . . . . . . . . . .0100101101$. | XX1 | I | 492 | v3.0 |  |  | \|xv|l | Load VSX Vector Left-justified with Length |
|  | XX1 | 1 | 498 | v3.0 |  |  | \|xwwsx | Load VSX Vector Word \& Splat Indexed |
|  | XX1 | 1 | 493 | v3.0 |  |  | Ixvx | Load VSX Vector Indexed |
| $000100 \ldots \ldots . . . .1 . . . . . . . .110000$ | VA | 1 | 81 | v3.0 |  |  | maddhd | Multiply-Add High Dword |
|  | VA | I | 81 | v3.0 |  |  | maddhdu | Multiply-Add High Dword Unsigned |
| $000100 \ldots \ldots . . . . . . . . . . . . . .110011$ | VA | 1 | 81 | v3.0 |  |  | maddld | Multiply-Add Low Dword |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 1 of 17)


Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 2 of 17)

| Instruction $^{\mathbf{1}}$ <br> 1 <br> $112345 \quad 67890 \quad 1211111112 \quad 22222 \quad 2222333$ <br> 67890 <br> $12345 \quad 678901$ |  | $\begin{aligned} & \text { ron } \\ & \text { on } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & \underset{\sim}{0} \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \frac{0}{n} \\ & \text { No } \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $000100 \ldots \ldots . . . . . . . . .11110001101$ | VX | \| | 347 | v3.0 |  |  | vextuwrx | Vector Extract Unsigned Word Right-Indexed |
| $000100 \ldots \ldots 1 \ldots \ldots 01100001101$ | VX | I | 270 | v3.0 |  |  | vinsertb | Vector Insert Byte |
| $000100 \ldots . .1 \ldots . . . . . .101111001101$ | VX | I | 270 | v3.0 |  |  | vinsertd | Vector Insert Dword |
| $000100 \ldots . .1 \ldots . . . . . .101101001101$ | VX | I | 270 | v3.0 |  |  | vinserth | Vector Insert Hword |
| $000100 \ldots . .1 \ldots \ldots .101110001101$ | VX | 1 | 270 | v3.0 |  |  | vinsertw | Vector Insert Word |
| $000100 \ldots \ldots .1111100000000001$ | VX | I | 357 | v3.0 |  |  | vmul10cuq | Vector Multiply-by-10 \& write Carry Unsigned Qword |
| $000100 \ldots \ldots . \ldots \ldots .00001000001$ | VX | I | 357 | v3.0 |  |  | vmul10ecuq | Vector Multiply-by-10 Extended \& write Carry Unsigned Qword |
|  | VX | I | 357 | v3.0 |  |  | vmul10euq | Vector Multiply-by-10 Extended Unsigned Qword |
| 000100 $\ldots \ldots \ldots .1111101000000001$ | VX | I | 357 | v3.0 |  |  | vmul10uq | Vector Multiply-by-10 Unsigned Qword |
| $000100 \ldots 0.10111 \ldots . .11000000010$ | VX | I | 295 | v3.0 |  |  | vnegd | Vector Negate Dword |
| $000100 \ldots . .00110 \ldots \ldots 11000000010$ | VX | 1 | 295 | v3.0 |  |  | vnegw | Vector Negate Word |
| $000100 \ldots \ldots . . . . . . . . . . . . . . .111011$ | VA | 1 | 262 | v3.0 |  |  | vpermr | Vector Permute Right-indexed |
| $000100 \ldots . .01001 \ldots \ldots 11000000010$ | VX | I | 317 | v3.0 |  |  | vprtybd | Vector Parity Byte Dword |
| $000100 \ldots 0.1010 \ldots \ldots 11000000010$ | VX | I | 317 | v3.0 |  |  | vprtybq | Vector Parity Byte Qword |
|  | VX | I | 317 | v3.0 |  |  | vprtybw | Vector Parity Byte Word |
| $000100 \ldots \ldots . . . . . . . . .100011000101$ | VX | I | 323 | v3.0 |  |  | vrldmi | Vector Rotate Left Dword then Mask Insert |
| $000100 \ldots \ldots 00111000101$ | VX | I | 323 | v3.0 |  |  | vrldnm | Vector Rotate Left Dword then AND with Mask |
| $000100 \ldots \ldots . \ldots \ldots .00010000101$ | VX | 1 | 322 | v3.0 |  |  | vrlwmi | Vector Rotate Left Word then Mask Insert |
| $000100 \ldots \ldots . \ldots \ldots . \ldots 0110000101$ | VX | I | 322 | v3.0 |  |  | vrlwnm | Vector Rotate Left Word then AND with Mask |
| $000100 \ldots \ldots . \ldots \ldots 11101000100$ | X | I | 267 | v3.0 |  |  | vslv | Vector Shift Left Variable |
| $000100 \ldots \ldots . . . . . . . .11100000100$ | X | I | 267 | v3.0 |  |  | VSTV | Vector Shift Right Variable |
| $011111111 . .111111111100000111101$ | X | 11 | 880 | v3.0 |  |  | wait | Wait for Interrupt |
| 111111.... $00000 \ldots 11001001001$ | X | I | 513 | v3.0 |  |  | xsabsqp | VSX Scalar Absolute QP |
| 111111............ 0000000100. | X | I | 521 | v3.0 |  |  | xsaddqp[0] | VSX Scalar Add QP |
| 111100 $\ldots \ldots . \ldots . . . . . .00000011 \ldots$ | XX3 | I | 525 | v3.0 |  |  | xscmpeqdp | VSX Scalar Compare Equal Double-Precision |
| 111100 $\ldots .11 \ldots . . . . . . .00111011 . .1$ | XX3 | I | 523 | v3.0 |  |  | xscmpexpdp | VSX Scalar Compare Exponents DP |
| 111111...11 .... ..... 00101001001 | X | 1 | 524 | v3.0 |  |  | xscmpexpqp | VSX Scalar Compare Exponents QP |
| 111100 .... .... .... 00010 011... | XX3 | I | 526 | v3.0 |  |  | xscmpgedp | VSX Scalar Compare Greater Than or Equal Double-Precision |
| 111100 $\ldots \ldots \ldots . \ldots \ldots$ | XX3 | I | 527 | v3.0 |  |  | xscmpgtdp | VSX Scalar Compare Greater Than Double-Precision |
|  | XX3 | I | 528 | v3.0 |  |  | xscmpnedp | VSX Scalar Compare Not Equal Double-Precision |
| 111111...11 ......... 00100001001 | X | I | 531 | v3.0 |  |  | xscmpoqp | VSX Scalar Compare Ordered QP |
| 1111111..11 $\ldots \ldots \ldots 10100001001$ | X | I | 534 | v3.0 |  |  | xscmpuqp | VSX Scalar Compare Unordered QP |
| 111111............ 00011001001 | X | 1 | 535 | v3.0 |  |  | xscpsgnqp | VSX Scalar Copy Sign QP |
| 111100 $\ldots . .10001 \ldots . .101011011 .$. | XX2 | I | 536 | v3.0 |  |  | xscvdphp | VSX Scalar Convert DP to HP |
| 111111.... $10110 \ldots . .11010001001$ | X | I | 537 | v3.0 |  |  | xscvdpqp | VSX Scalar Convert DP to QP |
| 111100 $\ldots . .10000 \ldots . .101011011$. | XX2 | I | 548 | v3.0 |  |  | xscvhpdp | VSX Scalar Convert HP to DP |
| 111111.... $10100 \ldots \ldots 1101000100$. | X | I | 549 | v3.0 |  |  | xscvqpdp[0] | VSX Scalar Convert QP to DP |
| 111111.... 11001.... 11010 001001 | X | I | 550 | v3.0 |  |  | xscvqpsdz | VSX Scalar Convert QP to Signed Dword truncate |
| 111111.... $01001 \ldots \ldots 11010001001$ | X | I | 552 | v3.0 |  |  | xscvqpswz | VSX Scalar Convert QP to Signed Word truncate |
| 111111.... 10001.... 11010 001001 | X | I | 554 | v3.0 |  |  | xscvqpudz | VSX Scalar Convert QP to Unsigned Dword truncate |
| 111111.... $00001 \ldots \ldots 11010001001$ | X | 1 | 556 | v3.0 |  |  | xscvqpuwz | VSX Scalar Convert QP to Unsigned Word truncate |
| 111111.... $01010 \ldots . .11010001001$ | X | I | 558 | v3.0 |  |  | xscvsdqp | VSX Scalar Convert Signed Dword to QP |
| 111111.... $00010 \ldots \ldots 11010001001$ | X | I | 562 | v3.0 |  |  | xscvudqp | VSX Scalar Convert Unsigned Dword to QP |
| 111111.............. 1000100100. | X | I | 566 | v3.0 |  |  | xsdivqp[0] | VSX Scalar Divide QP |
| 111100 $\ldots . . . . . . . . . . .1110010110$. | XX1 | I | 570 | v3.0 |  |  | xsiexpdp | VSX Scalar Insert Exponent DP |
| 111111............. 11011001001 | X | 1 | 571 | v3.0 |  |  | xsiexpqp | VSX Scalar Insert Exponent QP |
| 111111......... $\ldots . . .0110000100$. | X | I | 578 | v3.0 |  |  | xsmaddqp[0] | VSX Scalar Multiply-Add QP |
| 111100 $\ldots \ldots \ldots \ldots \ldots . . . .10000000 \ldots$ | XX3 | I | 583 | v3.0 |  |  | xsmaxcdp | VSX Scalar Maximum Type-C Double-Precision |
| 111100 $\ldots \ldots \ldots \ldots \ldots 10010000 \ldots$ | XX3 | I | 585 | v3.0 |  |  | xsmaxjdp | VSX Scalar Maximum Type-J Double-Precision |
| 111100 $\ldots \ldots \ldots \ldots 10001000 \ldots$ | XX3 | 1 | 589 | v3.0 |  |  | xsmincdp | VSX Scalar Minimum Type-C Double-Precision |
| 111100 $\ldots \ldots \ldots \ldots \ldots . . .10011000 \ldots$ | XX3 | I | 591 | v3.0 |  |  | xsminjdp | VSX Scalar Minimum Type-J Double-Precision |
| 111111.... $\ldots . . . . . . .0110100100$. | X | I | 599 | v3.0 |  |  | xsmsubqp[0] | VSX Scalar Multiply-Subtract QP |
| 111111............ 00000100100. | X | I | 604 | v3.0 |  |  | xsmulqp[0] | VSX Scalar Multiply QP |
| 111111.... $01000 \ldots . .11001001001$ | X | I | 608 | v3.0 |  |  | xsnabsqp | VSX Scalar Negative Absolute QP |
| 1111111.... 10000.... 11001001001 | X | I | 609 | v3.0 |  |  | xsnegqp | VSX Scalar Negate QP |
| 111111............... 0111000100. | X | 1 | 618 | v3.0 |  |  | xsnmaddqp[0] | VSX Scalar Negative Multiply-Add QP |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 3 of 17)


Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 4 of 17)

| Instruction $^{\mathbf{1}}$ 1 $111111 \quad 11112 \quad 22222 \quad 222233$ $012345678901234567890 \quad 12345678901$ |  | $\begin{aligned} & \text { ro } \\ & \text { on } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & \underset{\sim}{0} \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \frac{0}{n} \\ & \text { No } \\ & \hline \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $011111 \ldots . . . . . . . . . . . .0010001100$. | XX1 | I | 501 | v2.07 |  |  | stxsiwx | Store VSX Scalar as Integer Word Indexed |
| 011111..... .... ..... 1010001100. | XX1 | I | 503 | v2.07 |  |  | stxsspx | Store VSX Scalar SP Indexed |
| $01111111111 \ldots . . .1111111100011101$ | X | 11 | 895 | v2.07 |  |  | tabort. | Transaction Abort \& record |
| 011111 $\ldots \ldots . . . . . . . . . . .11001011101$ | X | 11 | 897 | v2.07 |  |  | tabortdc. | Transaction Abort Dword Conditional \& record |
|  | X | 11 | 897 | v2.07 |  |  | tabortdci. | Transaction Abort Dword Conditional Immediate \& record |
| $011111 \ldots \ldots . . . . . . . . . .11000011101$ | X | I\| | 896 | v2.07 |  |  | tabortwc. | Transaction Abort Word Conditional \& record |
|  | X | 11 | 896 | v2.07 |  |  | tabortwci. | Transaction Abort Word Conditional Immediate \& record |
| 011111.111. 11111 \|111110100 011101 | X | 11 | 893 | v2.07 |  |  | tbegin. | Transaction Begin \& record |
| 011111 ...11 11111 11111 10110 011101 | X | 11 | 898 | v2.07 |  |  | tcheck | Transaction Check \& record |
| 011111.1111 11111 1111\| 10101 011101 | X | 11 | 894 | v2.07 |  |  | tend. | Transaction End \& record |
| 01111111111111111111111111011101 | X | 111 | 970 | v2.07 | P |  | trechkpt. | Transaction Recheckpoint \& record |
| 011111 \||11| $\ldots . . .1\|1\| \mid 11101011101$ | X | 111 | 969 | v2.07 | P |  | treclaim. | Transaction Reclaim \& record |
| 0111111111.111111111110111011101 | X | 11 | 898 | v2.07 |  |  | tsr. | Transaction Suspend or Resume \& record |
| $000100 \ldots \ldots . \ldots . . . . .00101000000$ | VX | I | 275 | v2.07 |  |  | vaddcuq | Vector Add \& write Carry Unsigned Qword |
| 000100 $\ldots \ldots . \ldots . . . . . . . . . . . . .111101$ | VA | 1 | 275 | v2.07 |  |  | vaddecuq | Vector Add Extended \& write Carry Unsigned Qword |
| 000100 $\ldots \ldots . . . . . . . . . . . . . . . . .111100$ | VA | 1 | 275 | v2.07 |  |  | vaddeuqm | Vector Add Extended Unsigned Qword Modulo |
| $000100 \ldots \ldots . \ldots \ldots 0001100000$ | VX | 1 | 272 | v2.07 |  |  | vaddudm | Vector Add Unsigned Dword Modulo |
| $000100 \ldots \ldots . \ldots \ldots .00100000000$ | VX | I | 272 | v2.07 |  |  | vadduqm | Vector Add Unsigned Qword Modulo |
| $000100 \ldots \ldots . \ldots \ldots .10101001100$ | VX | 1 | 349 | v2.07 |  |  | vbpermq | Vector Bit Permute Qword |
| $000100 \ldots \ldots . \ldots \ldots .10100001000$ | VX | I | 336 | v2.07 |  |  | vcipher | Vector AES Cipher |
| $000100 \ldots \ldots \ldots \ldots 10100001001$ | VX | 1 | 336 | v2.07 |  |  | vcipherlast | Vector AES Cipher Last |
| $000100 \ldots \ldots 11111 \ldots \ldots .11100000010$ | VX | I | 343 | v2.07 |  |  | vclzb | Vector Count Leading Zeros Byte |
| $000100 \ldots . .11111 \ldots \ldots .11111000010$ | VX | 1 | 343 | v2.07 |  |  | vclzd | Vector Count Leading Zeros Dword |
| $000100 \ldots . .11111 \ldots . .11101000010$ | VX | I | 343 | v2.07 |  |  | vclzh | Vector Count Leading Zeros Hword |
| $000100 \ldots . .11111 \ldots \ldots 1111000010$ | VX | 1 | 343 | v2.07 |  |  | vclzw | Vector Count Leading Zeros Word |
| 000100 $\ldots \ldots . \ldots . . . . . . .0011000111$ | VC | I | 307 | v2.07 |  |  | vcmpequd[.] | Vector Compare Equal Unsigned Dword |
| $000100 \ldots \ldots . \ldots \ldots .1111000111$ | VC | I | 308 | v2.07 |  |  | vcmpgtsd[.] | Vector Compare Greater Than Signed Dword |
| $000100 \ldots \ldots . \ldots . . . . . .1011000111$ | VC | 1 | 310 | v2.07 |  |  | vcmpgtud[.] | Vector Compare Greater Than Unsigned Dword |
| $000100 \ldots \ldots . \ldots \ldots .11010000100$ | VX | 1 | 315 | v2.07 |  |  | veqv | Vector Logical Equivalence |
| 000100 $\ldots \ldots 11111 \ldots \ldots 10100001100$ | VX | 1 | 342 | v2.07 |  |  | vgbbd | Vector Gather Bits by Byte by Dword |
| $000100 \ldots \ldots . . . . . . . . . .00111000010$ | VX | 1 | 302 | v2.07 |  |  | vmaxsd | Vector Maximum Signed Dword |
| $000100 \ldots \ldots . \ldots \ldots, \ldots 0011000010$ | VX | I | 302 | v2.07 |  |  | vmaxud | Vector Maximum Unsigned Dword |
| $000100 \ldots \ldots . \ldots \ldots .01111000010$ | VX | I | 304 | v2.07 |  |  | vminsd | Vector Minimum Signed Dword |
| 000100 $\ldots \ldots . \ldots \ldots . . . . .01011000010$ | VX | 1 | 304 | v2.07 |  |  | vminud | Vector Minimum Unsigned Dword |
| $000100 \ldots \ldots . \ldots \ldots 11110001100$ | VX | I | 259 | v2.07 |  |  | vmrgew | Vector Merge Even Word |
| $000100 \ldots \ldots . \ldots \ldots \ldots 11010001100$ | VX | I | 259 | v2.07 |  |  | vmrgow | Vector Merge Odd Word |
|  | VX | I | 285 | v2.07 |  |  | vmulesw | Vector Multiply Even Signed Word |
| 000100 $\ldots \ldots \ldots \ldots \ldots . \ldots 1010001000$ | VX | I | 285 | v2.07 |  |  | vmuleuw | Vector Multiply Even Unsigned Word |
| $000100 \ldots \ldots . \ldots \ldots 0 . \ldots 0110001000$ | VX | 1 | 285 | v2.07 |  |  | vmulosw | Vector Multiply Odd Signed Word |
| $000100 \ldots \ldots . \ldots \ldots 0.10010001000$ | VX | I | 285 | v2.07 |  |  | vmulouw | Vector Multiply Odd Unsigned Word |
| $000100 \ldots \ldots . \ldots \ldots 0.10010001001$ | VX | I | 286 | v2.07 |  |  | vmuluwm | Vector Multiply Unsigned Word Modulo |
| 000100 $\ldots \ldots \ldots \ldots \ldots .10110000100$ | VX | I | 315 | v2.07 |  |  | vnand | Vector Logical NAND |
| $000100 \ldots \ldots . \ldots \ldots .10101001000$ | VX | I | 337 | v2.07 |  |  | vncipher | Vector AES Inverse Cipher |
| $000100 \ldots \ldots . \ldots \ldots . \ldots 10101001001$ | VX | I | 337 | v2.07 |  |  | vncipherlast | Vector AES Inverse Cipher Last |
| 000100 $\ldots \ldots \ldots \ldots \ldots . . . . . .10101000100$ | VX | I | 316 | v2.07 |  |  | vorc | Vector Logical OR with Complement |
| 000100 $\ldots \ldots . \ldots . . . . . . . . . . . .101101$ | VA | I | 341 | v2.07 |  |  | vpermxor | Vector Permute \& Exclusive-OR |
|  | VX | I | 250 | v2.07 |  |  | vpksdss | Vector Pack Signed Dword Signed Saturate |
| $000100 \ldots \ldots . \ldots \ldots . . . . .10101001110$ | VX | 1 | 251 | v2.07 |  |  | vpksdus | Vector Pack Signed Dword Unsigned Saturate |
| $000100 \ldots \ldots . \ldots \ldots .10001001110$ | VX | I | 253 | v2.07 |  |  | vpkudum | Vector Pack Unsigned Dword Unsigned Modulo |
| $000100 \ldots \ldots . \ldots \ldots 10011001110$ | VX | I | 253 | v2.07 |  |  | vpkudus | Vector Pack Unsigned Dword Unsigned Saturate |
| $000100 \ldots \ldots . \ldots \ldots .10000001000$ | VX | I | 339 | v2.07 |  |  | vpmsumb | Vector Polynomial Multiply-Sum Byte |
| $000100 \ldots \ldots . \ldots \ldots .10011001000$ | VX | I | 339 | v2.07 |  |  | vpmsumd | Vector Polynomial Multiply-Sum Dword |
| 000100 $\ldots \ldots \ldots \ldots . . . . .10001001000$ | VX | 1 | 340 | v2.07 |  |  | vpmsumh | Vector Polynomial Multiply-Sum Hword |
| 000100 $\ldots \ldots \ldots \ldots \ldots .10010001000$ | VX | I | 340 | v2.07 |  |  | vpmsumw | Vector Polynomial Multiply-Sum Word |
| $000100 \ldots . .11111 \ldots \ldots .11100000011$ | VX | 1 | 348 | v2.07 |  |  | vpopentb | Vector Population Count Byte |
| $000100 \ldots . .11111 \ldots . .11111000011$ | VX | 1 | 348 | v2.07 |  |  | vpopentd | Vector Population Count Dword |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 5 of 17)

| Instruction $^{\mathbf{1}}$ 1 $11111111112 \quad 22222222233$ $01234567890 \quad 123456789012345678901$ |  | $\begin{aligned} & \text { ro } \\ & 0 \\ & 0 \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & 0 \\ & 0 \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \frac{0}{n} \\ & \frac{0}{0} \\ & \hline \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $000100 \ldots . .11111 \ldots . .11101000011$ | VX | I | 348 | v2.07 |  |  | vpopenth | Vector Population Count Hword |
| $000100 \ldots . .11111 \ldots . .11110000011$ | VX | I | 348 | v2.07 |  |  | vpopentw | Vector Population Count Word |
|  | VX | I | 318 | v2.07 |  |  | vrld | Vector Rotate Left Dword |
| $000100 \ldots \ldots . \ldots 1111110111001000$ | VX | 1 | 337 | v2.07 |  |  | vsbox | Vector AES SubBytes |
| 000100 $\ldots \ldots . \ldots \ldots \ldots 11011000010$ | VX | I | 338 | v2.07 |  |  | vshasigmad | Vector SHA-512 Sigma Dword |
| $000100 \ldots \ldots . . . . . . . . .11010000010$ | VX | I | 338 | v2.07 |  |  | vshasigmaw | Vector SHA-256 Sigma Word |
| $000100 \ldots \ldots . . . . . . . . .10111000100$ | VX | I | 319 | v2.07 |  |  | vsld | Vector Shift Left Dword |
| 000100 $\ldots \ldots . \ldots \ldots . . . . .01111000100$ | VX | I | 321 | v2.07 |  |  | vsrad | Vector Shift Right Algebraic Dword |
| 000100 $\ldots \ldots . . . . . . . . .11011000100$ | VX | I | 320 | v2.07 |  |  | vsrd | Vector Shift Right Dword |
| $000100 \ldots \ldots . \ldots \ldots 10101000000$ | VX | I | 281 | v2.07 |  |  | vsubcuq | Vector Subtract \& write Carry Unsigned Qword |
| 000100 $\ldots \ldots . . . . . . . . . . . . . . . . .111111$ | VA | I | 281 | v2.07 |  |  | vsubecuq | Vector Subtract Extended \& write Carry Unsigned Qword |
| $000100 \ldots \ldots . . . . . . . . . . . . . . .111110$ | VA | 1 | 281 | v2.07 |  |  | vsubeuqm | Vector Subtract Extended Unsigned Qword Modulo |
| $000100 \ldots \ldots . \ldots \ldots 10011000000$ | VX | I | 279 | v2.07 |  |  | vsubudm | Vector Subtract Unsigned Dword Modulo |
| $000100 \ldots \ldots . \ldots \ldots 10100000000$ | VX | 1 | 281 | v2.07 |  |  | vsubuqm | Vector Subtract Unsigned Qword Modulo |
| 000100 $\ldots . .111111 \ldots . .11001001110$ | VX | 1 | 256 | v2.07 |  |  | vupkhsw | Vector Unpack High Signed Word |
| 000100 $\ldots . .111111 \ldots . .11011001110$ | VX | 1 | 256 | v2.07 |  |  | vupklsw | Vector Unpack Low Signed Word |
| 111100 $\ldots \ldots . \ldots \ldots . \ldots 0000000 \ldots$ | XX3 | 1 | 519 | v2.07 |  |  | xsaddsp | VSX Scalar Add SP |
| 111100 $\ldots . .111111 \ldots . .100001011 .$. | XX2 | 1 | 539 | v2.07 |  |  | xscvdpspn | VSX Scalar Convert DP to SP Non-signalling |
| 111100 $\ldots . . .11111 \ldots \ldots 101001011 .$. | XX2 | 1 | 560 | v2.07 |  |  | xscvspdpn | VSX Scalar Convert SP to DP Non-signalling |
| 111100 $\ldots . .11111 \ldots . . .100111000 .$. | XX2 | 1 | 561 | v2.07 |  |  | xscvsxdsp | VSX Scalar Convert Signed Dword to SP |
| 111100 $\ldots . .11111 \ldots \ldots 100101000 .$. | XX2 | 1 | 563 | v2.07 |  |  | xscvuxdsp | VSX Scalar Convert Unsigned Dword to SP |
| 111100 ............. $00011000 \ldots$ | XX3 | I | 568 | v2.07 |  |  | xsdivsp | VSX Scalar Divide SP |
|  | XX3 | I | 575 | v2.07 |  |  | xsmaddasp | VSX Scalar Multiply-Add Type-A SP |
|  | XX3 | 1 | 575 | v2.07 |  |  | xsmaddmsp | VSX Scalar Multiply-Add Type-M SP |
| 111100 $\ldots . . \ldots . . . . . .00010001 \ldots$ | XX3 | 1 | 596 | v2.07 |  |  | xsmsubasp | VSX Scalar Multiply-Subtract Type-A SP |
| 111100 $\ldots \ldots . \ldots . . . . . .00011001 \ldots$ | XX3 | 1 | 596 | v2.07 |  |  | xsmsubmsp | VSX Scalar Multiply-Subtract Type-M SP |
| 111100 $\ldots \ldots \ldots \ldots \ldots 00010000 \ldots$ | XX3 | 1 | 606 | v2.07 |  |  | xsmulsp | VSX Scalar Multiply SP |
| 111100 $\ldots \ldots \ldots \ldots \ldots 10000001 \ldots$ | XX3 | 1 | 615 | v2.07 |  |  | xsnmaddasp | VSX Scalar Negative Multiply-Add Type-A SP |
| 111100 $\ldots \ldots . \ldots \ldots \ldots .10001001 \ldots$ | XX3 | 1 | 615 | v2.07 |  |  | xsnmaddmsp | VSX Scalar Negative Multiply-Add Type-M SP |
| 111100 $\ldots .$. $\ldots .$. $\ldots .$. 10010 $001 \ldots$ | XX3 | I | 624 | v2.07 |  |  | xsnmsubasp | VSX Scalar Negative Multiply-Subtract Type-A SP |
| 1111100 $\ldots \ldots . . . . . . . . .100111001 \ldots$ | XX3 | I | 624 | v2.07 |  |  | xsnmsubmsp | VSX Scalar Negative Multiply-Subtract Type-M SP |
| 111100 $\ldots . .11111 \ldots \ldots 000011010 .$. | XX2 | 1 | 635 | v2.07 |  |  | xsresp | VSX Scalar Reciprocal Estimate SP |
| 111100 $\ldots . .11111 \ldots . .100011001 .$. | XX2 | 1 | 640 | v2.07 |  |  | xsrsp | VSX Scalar Round DP to SP |
| 111100 $\ldots . .11111 \ldots . . .000001010 .$. | XX2 | 1 | 642 | v2.07 |  |  | xsrsqrtesp | VSX Scalar Reciprocal Square Root Estimate SP |
| 111100 $\ldots . .11111 \ldots . . .000001011 .$. | XX2 | 1 | 646 | v2.07 |  |  | xssqrtsp | VSX Scalar Square Root SP |
| 111100 $\ldots \ldots . \ldots \ldots 00001000 \ldots$ | XX3 | 1 | 651 | v2.07 |  |  | xssubsp | VSX Scalar Subtract SP |
| 1111100 $\ldots \ldots . . . . . . . . . . .101111010 \ldots$ | XX3 | I | 772 | v2.07 |  |  | xxleqv | VSX Vector Logical Equivalence |
| 111100 $\ldots \ldots . . . . . . . . . .101101010 \ldots$ | XX3 | 1 | 772 | v2.07 |  |  | xxInand | VSX Vector Logical NAND |
| 111100 $\ldots \ldots . \ldots . . . . . . .10101010 \ldots$ | XX3 | 1 | 773 | v2.07 |  |  | xxlorc | VSX Vector Logical OR with Complement |
|  | XO | 1 | 110 | v2.06 |  |  | addg6s | Add \& Generate Sixes |
|  | X | I | 99 | v2.06 |  |  | bpermd | Bit Permute Dword |
| 011111......... 1111101001110101 | X | I | 110 | v2.06 |  |  | cbcdtd | Convert Binary Coded Decimal To Declets |
| 011111 $\ldots \ldots . . . . .1111101000110101$ | X | I | 110 | v2.06 |  |  | cdtbcd | Convert Declets To Binary Coded Decimal |
| 111011 $\ldots . . .11111 \ldots \ldots .1100100010$, | X | I | 217 | v2.06 |  |  | dcffix[.] | DFP Convert From Fixed |
| 011111............... . 110101001. | XO | 1 | 83 | v2.06 |  | SR | divde[0][.] | Divide Dword Extended |
| 011111 $\ldots .$. $\ldots .$. $\ldots .$. <br> 11100 01001.   | XO | 1 | 83 | v2.06 |  | SR | divdeu[0][.] | Divide Dword Extended Unsigned |
| 011111 $\ldots . . . . . . . . . . . ~$ 110101011. | XO | 1 | 77 | v2.06 |  | SR | divwe[0][.] | Divide Word Extended |
|  | XO | I | 77 | v2.06 |  | SR | divweu[0][.] | Divide Word Extended Unsigned |
| 111011 $\ldots . .111111 . . .11101001110$. | X | 1 | 165 | v2.06 |  |  | fcfids[.] | Floating Convert From Integer Dword Single |
| 111111.... $11111 \ldots \ldots .1111001110$. | X | I | 165 | v2.06 |  |  | fcfidu[.] | Floating Convert From Integer Dword Unsigned |
| 111011 $\ldots . . .111111 . . .11111001110$. | X | 1 | 166 | v2.06 |  |  | fcfidus[.] | Floating Convert From Integer Dword Unsigned Single |
| 111111 $\ldots . . .11111 \ldots . . .1110101110$. | X | I | 161 | v2.06 |  |  | fctidu[.] | Floating Convert To Integer Dword Unsigned |
| 1111111.... $11111 \ldots \ldots 1110101111$. | X | 1 | 162 | v2.06 |  |  | fctiduz[.] | Floating Convert To Integer Dword Unsigned truncate |
| 111111.... $11111 \ldots . . .0010001110$. | X | 1 | 163 | v2.06 |  |  | fctiwu[.] | Floating Convert To Integer Word Unsigned |
| 111111 $\ldots \ldots . .11111 \ldots . .0010001111$. | X | 1 | 164 | v2.06 |  |  | fctiwuz[.] | Floating Convert To Integer Word Unsigned truncate |
| 111111...11 ..... ..... 00100000001 | X | I | 157 | v2.06 |  |  | ftdiv | Floating Test for software Divide |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 6 of 17)

| Instruction $^{\mathbf{1}}$ 1 $111111 \quad 11112 \quad 22222 \quad 2222333$ $012345678901234567890 \quad 12345678901$ |  | $\begin{aligned} & \text { Yo } \\ & \text { ò } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & \underset{\sim}{0} \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \frac{0}{n} \\ & \text { No } \\ & \hline \end{aligned}$ |  | + 0 0 0 0 0 0 0 |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 111111...11 11111.... 0010100001 | X | 1 | 157 | v2.06 |  |  | ftsqrt | Floating Test for software Square Root |
| $011111 . . . .1 . . . . . . . .0000110100$. | X | 11 | 868 | v2.06 |  |  | Ibarx | Load Byte And Reserve Indexed |
| 011111.............. 10000101001 | X | I | 62 | v2.06 |  |  | Idbrx | Load Dword Byte-Reverse Indexed |
| 011111 $\ldots \ldots . . . . . . . . . . . .11011101111$ | X | I | 144 | v2.06 |  |  | lfiwzx | Load Floating as Integer Word \& Zero Indexed |
| $011111 \ldots . . . . . . . . . . . .0001110100$. | X | 11 | 869 | v2.06 |  |  | Iharx | Load Hword And Reserve Indexed Xform |
| $011111 \ldots \ldots . . . . . . . . . .1001001100$. | XX1 | 1 | 481 | v2.06 |  |  | Ixsdx | Load VSX Scalar Dword Indexed |
| 011111.............. 1101001100. | XX1 | I | 489 | v2.06 |  |  | Ixvd2x | Load VSX Vector Dword*2 Indexed |
| 011111 $\ldots .$. $\ldots . .$. $\ldots 1010$ 01100 | XX1 | I | 495 | v2.06 |  |  | Ixvdsx | Load VSX Vector Dword \& Splat Indexed |
| 011111.............. 1100001100. | XX1 | 1 | 497 | v2.06 |  |  | Ixvw4x | Load VSX Vector Word*4 Indexed |
| 011111......... 11111101111110101 | X | I | 98 | v2.06 |  |  | popentd | Population Count Dword |
| 011111......... 11111101011110101 | X | I | 96 | v2.06 |  |  | popentw | Population Count Words |
| 011111 $\ldots \ldots . . . . . . . . . . .10101101101$ | X | 11 | 870 | v2.06 |  |  | stbex. | Store Byte Conditional Indexed \& record |
| $011111 \ldots . . . . . . . . . . . .10100101001$ | X | I | 62 | v2.06 |  |  | stdbrx | Store Dword Byte-Reverse Indexed |
| $011111 \ldots . . . . . . . . . . . .10110101101$ | X | 11 | 871 | v2.06 |  |  | sthcx. | Store Hword Conditional Indexed \& record |
|  | XX1 | I | 499 | v2.06 |  |  | stxsdx | Store VSX Scalar Dword Indexed |
| 11111 $\ldots . .$. $\ldots . .$. 11110 01100 | XX1 | 1 | 505 | v2.06 |  |  | stxvd2x | Store VSX Vector Dword*2 Indexed |
| $011111 \ldots \ldots . . . . . . . . . .1110001100$. | XX1 | I | 507 | v2.06 |  |  | stxvw4x | Store VSX Vector Word*4 Indexed |
| 111100 $\ldots . .11111 \ldots \ldots .101011001$. | XX2 | 1 | 513 | v2.06 |  |  | xsabsdp | VSX Scalar Absolute DP |
|  | XX3 | 1 | 514 | v2.06 |  |  | xsadddp | VSX Scalar Add DP |
| 111100 $\ldots .11 \ldots \ldots . . . . .100101011 . .1$ | XX3 | 1 | 529 | v2.06 |  |  | xscmpodp | VSX Scalar Compare Ordered DP |
| 111100 $\ldots .11 \ldots \ldots . . . . .100100011 . .1$ | XX3 | 1 | 532 | V2.06 |  |  | xscmpudp | VSX Scalar Compare Unordered DP |
| 111100 $\ldots \ldots \ldots . . . . . . .10110000 \ldots$ | XX3 | \| | 535 | v2.06 |  |  | xscpsgndp | VSX Scalar Copy Sign DP |
| 111100 $\ldots \ldots .11111 \ldots \ldots .100001001$. | XX2 | I | 538 | v2.06 |  |  | xscvdpsp | VSX Scalar Convert DP to SP |
| 111100 $\ldots . .11111 \ldots . .101011000$. | XX2 | 1 | 539 | v2.06 |  |  | xscvdpsxds | VSX Scalar Convert DP to Signed Dword truncate |
| 111100 $\ldots . .11111 \ldots . .001011000$. | XX2 | 1 | 542 | v2.06 |  |  | xscvdpsxws | VSX Scalar Convert DP to Signed Word truncate |
| 111100 $\ldots \ldots .11111 \ldots \ldots 101001000$. | XX2 | I | 544 | v2.06 |  |  | xscvdpuxds | VSX Scalar Convert DP to Unsigned Dword truncate |
| 111100 $\ldots . .11111 \ldots . .001001000$. | XX2 | I | 546 | v2.06 |  |  | xscvdpuxws | VSX Scalar Convert DP to Unsigned Word truncate |
| 111100 $\ldots . .11111 \ldots . .101001001$. | XX2 | 1 | 559 | v2.06 |  |  | xscvspdp | VSX Scalar Convert SP to DP |
| 111100 $\ldots . .111111 \ldots . .101111000$ | XX2 | 1 | 561 | v2.06 |  |  | xscvsxddp | VSX Scalar Convert Signed Dword to DP |
| 111100 $\ldots . .11111 \ldots \ldots 101101000$. | XX2 | 1 | 563 | v2.06 |  |  | xscvuxddp | VSX Scalar Convert Unsigned Dword to DP |
| 111100 $\ldots \ldots . . . . . . . . . . ~ 00111000 \ldots$ | XX3 | I | 564 | v2.06 |  |  | xsdivdp | VSX Scalar Divide DP |
| 111100 $\ldots \ldots . \ldots \ldots . \ldots .00100001 \ldots$ | XX3 | I | 572 | v2.06 |  |  | xsmaddadp | VSX Scalar Multiply-Add Type-A DP |
| 111100 $\ldots \ldots . \ldots \ldots . \ldots 00101001 \ldots$ | XX3 | 1 | 572 | v2.06 |  |  | xsmaddmdp | VSX Scalar Multiply-Add Type-M DP |
| 111100 $\ldots \ldots . . . . . . . . . .10100000 \ldots$ | XX3 | 1 | 581 | v2.06 |  |  | xsmaxdp | VSX Scalar Maximum DP |
| 111100 $\ldots \ldots \ldots \ldots \ldots 10101000 \ldots$ | XX3 | 1 | 587 | v2.06 |  |  | xsmindp | VSX Scalar Minimum DP |
| 111100 $\ldots \ldots . \ldots \ldots . \ldots 00110001 \ldots$ | XX3 | I | 593 | v2.06 |  |  | xsmsubadp | VSX Scalar Multiply-Subtract Type-A DP |
| 111100 $\ldots \ldots . \ldots . . . . . . . .00111001 \ldots$ | XX3 | I | 593 | v2.06 |  |  | xsmsubmdp | VSX Scalar Multiply-Subtract Type-M DP |
| 111100 $\ldots \ldots . . . . . . . . . . . ~ 00110000 . ~$ | XX3 | 1 | 602 | v2.06 |  |  | xsmuldp | VSX Scalar Multiply DP |
| 111100 $\ldots \ldots .11111 \ldots \ldots 101101001$. | XX2 | 1 | 608 | v2.06 |  |  | xsnabsdp | VSX Scalar Negative Absolute DP |
| 111100 $\ldots . .111111 \ldots . .101111001$. | XX2 | I | 609 | v2.06 |  |  | xsnegdp | VSX Scalar Negate DP |
| 111100 $\ldots \ldots \ldots \ldots \ldots .10100001 \ldots$ | XX3 | I | 610 | v2.06 |  |  | xsnmaddadp | VSX Scalar Negative Multiply-Add Type-A DP |
| 111100 $\ldots \ldots \ldots . . . . . . .10101001 \ldots$ | XX3 | I | 610 | v2.06 |  |  | xsnmaddmdp | VSX Scalar Negative Multiply-Add Type-M DP |
| 111100 $\ldots \ldots . \ldots \ldots \ldots .10110$ 001... | XX3 | । | 621 | v2.06 |  |  | xsnmsubadp | VSX Scalar Negative Multiply-Subtract Type-A DP |
| 111100 $\ldots \ldots . . . . . . . . . . .10111001 . .1$ | XX3 | I | 621 | v2.06 |  |  | xsnmsubmdp | VSX Scalar Negative Multiply-Subtract Type-M DP |
| 111100 $\ldots . .11111 \ldots . .001001001 .$. | XX2 | I | 630 | v2.06 |  |  | xsrdpi | VSX Scalar Round DP to Integral to Nearest Away |
| 111100 $\ldots . .11111 \ldots . . .001101011$. | XX2 | 1 | 631 | v2.06 |  |  | xsrdpic | VSX Scalar Round DP to Integral using Current rounding mode |
| 111100 $\ldots . . .11111 \ldots . . .001111001$. | XX2 | 1 | 632 | v2.06 |  |  | xsrdpim | VSX Scalar Round DP to Integral toward -Infinity |
| 111100 $\ldots . .11111 \ldots . . .001101001$. | XX2 | I | 632 | v2.06 |  |  | xsrdpip | VSX Scalar Round DP to Integral toward +Infinity |
| 111100 $\ldots . .11111 \ldots . .001011001$. | XX2 | 1 | 633 | v2.06 |  |  | xsrdpiz | VSX Scalar Round DP to Integral toward Zero |
| $111100 \ldots . .11111 \ldots . .001011010$. | XX2 | I | 634 | v2.06 |  |  | xsredp | VSX Scalar Reciprocal Estimate DP |
| 111100 $\ldots . .11111 \ldots \ldots 001001010$. | XX2 | I | 641 | v2.06 |  |  | xsrsqriedp | VSX Scalar Reciprocal Square Root Estimate DP |
| 111100 $\ldots . .11111 \ldots . . .001001011$. | XX2 | I | 643 | v2.06 |  |  | xssqrtdp | VSX Scalar Square Root DP |
| 111100 $\ldots \ldots \ldots . . . . . . .00101000 \ldots$ | XX3 | 1 | 647 | v2.06 |  |  | xssubdp | VSX Scalar Subtract DP |
| 111100 $\ldots .11 \ldots . . .1 . . .00111101 . .1$ | XX3 | I | 653 | v2.06 |  |  | xstdivdp | VSX Scalar Test for software Divide DP |
| 111100 $\ldots 1111111 \ldots \ldots 001101010.1$ | XX2 | 1 | 654 | v2.06 |  |  | xstsqrtdp | VSX Scalar Test for software Square Root DP |
| 111100..... $11111 \ldots . . .111011001$. | XX2 | 1 | 660 | v2.06 |  |  | xvabsdp | VSX Vector Absolute DP |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 7 of 17)

|  |  | $\begin{aligned} & \text { Yo } \\ & \text { ò } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & \tilde{\pi} \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \stackrel{0}{n} \\ & \frac{1}{\omega} \\ & \hline \end{aligned}$ |  |  | 0 <br> 1 <br> 0 <br> 1 <br> 1 | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $111100 \ldots . .11111 \ldots . . .110011001 .$. | XX2 | I | 660 | v2.06 |  |  | xvabssp | VSX Vector Absolute SP |
| 111100 ......... .... $01100000 \ldots$ | XX3 | 1 | 661 | v2.06 |  |  | xvadddp | VSX Vector Add DP |
| $111100 \ldots . . . . . . . . . . .01000000 \ldots$ | XX3 | 1 | 665 | v2.06 |  |  | xvaddsp | VSX Vector Add SP |
|  | XX3 | I | 667 | v2.06 |  |  | xvcmpeqdp[.] | VSX Vector Compare Equal DP |
| 111100 .............. 1000 011... | XX3 | 1 | 668 | v2.06 |  |  | xvcmpeqsp[.] | VSX Vector Compare Equal SP |
| 111100 $\ldots . . . . . . . . . . .1110011 \ldots$ | XX3 | 1 | 669 | v2.06 |  |  | xvcmpgedp[.] | VSX Vector Compare Greater Than or Equal DP |
| 111100 $\ldots . . . \ldots . . . . . . . .1010 ~ 011 . . . ~$ | XX3 | 1 | 670 | v2.06 |  |  | xvcmpgesp[.] | VSX Vector Compare Greater Than or Equal SP |
| 111100 ......... ..... .1101 011... | XX3 | I | 671 | v2.06 |  |  | xvcmpgtdp[.] | VSX Vector Compare Greater Than DP |
| $111100 \ldots . . . . . . . . . .1001011 \ldots$ | XX3 | 1 | 672 | v2.06 |  |  | xvcmpgtsp[.] | VSX Vector Compare Greater Than SP |
| $111100 \ldots . . . . . . . . .11110000 \ldots$ | XX3 | 1 | 675 | v2.06 |  |  | xvcpsgndp | VSX Vector Copy Sign DP |
| $111100 \ldots . . . . . . . . .111010000 \ldots$ | XX3 | 1 | 675 | v2.06 |  |  | xvcpsgnsp | VSX Vector Copy Sign SP |
| $111100 \ldots 11111 \ldots . . .110001001 .$. | XX2 | 1 | 676 | v2.06 |  |  | xvcvdpsp | VSX Vector Convert DP to SP |
| $111100 \ldots . .11111 \ldots . .111011000$. | XX2 | 1 | 677 | v2.06 |  |  | xvcvdpsxds | VSX Vector Convert DP to Signed Dword truncate |
| $111100 \ldots . .11111 \ldots . .1011011000$. | XX2 | I | 679 | v2.06 |  |  | xvcvdpsxws | VSX Vector Convert DP to Signed Word truncate |
| 111100 $\ldots . .111111 \ldots . .111001000$. | XX2 | I | 681 | v2.06 |  |  | xvcvdpuxds | VSX Vector Convert DP to Unsigned Dword truncate |
| $111100 \ldots .11111 \ldots . .011001000$. | XX2 | 1 | 683 | v2.06 |  |  | xvcvdpuxws | VSX Vector Convert DP to Unsigned Word truncate |
| $111100 \ldots 11111 \ldots 111001001 .$. | XX2 | 1 | 686 | v2.06 |  |  | xvcvspdp | VSX Vector Convert SP to DP |
| $111100 \ldots . .11111 \ldots . .110011000$. | XX2 | 1 | 688 | v2.06 |  |  | xvcvspsxds | VSX Vector Convert SP to Signed Dword truncate |
| $111100 \ldots 11111 \ldots . . .1010011000 .$. | XX2 | I | 690 | v2.06 |  |  | xvcvspsxws | VSX Vector Convert SP to Signed Word truncate |
| $111100 \ldots . .11111 \ldots 11000$ 1000... | XX2 | I | 692 | v2.06 |  |  | xvcvspuxds | VSX Vector Convert SP to Unsigned Dword truncate |
| $111100 \ldots 11111 \ldots . . .1010001000 .$. | XX2 | I | 694 | v2.06 |  |  | xvcvspuxws | VSX Vector Convert SP to Unsigned Word truncate |
| $111100 \ldots . .11111 \ldots . .111111000$. | XX2 | 1 | 696 | v2.06 |  |  | xvcvsxddp | VSX Vector Convert Signed Dword to DP |
| $111100 \ldots . .11111 \ldots . .110111000$. | XX2 | 1 | 696 | v2.06 |  |  | xvcvsxdsp | VSX Vector Convert Signed Dword to SP |
| $111100 \ldots 11111 \ldots 011111000 .$. | XX2 | 1 | 697 | v2.06 |  |  | xvCvSxwdp | VSX Vector Convert Signed Word to DP |
| 111100 $\ldots . .11111 \ldots . .1010111000 .$. | XX2 | 1 | 697 | v2.06 |  |  | xvcvsxwsp | VSX Vector Convert Signed Word to SP |
| 111100 $\ldots . . .11111 \ldots . .111101000$. | XX2 | 1 | 698 | v2.06 |  |  | xvcvuxddp | VSX Vector Convert Unsigned Dword to DP |
| 111100 $\ldots . .11111 \ldots . . .110101000$. | XX2 | + | 698 | v2.06 |  |  | xvcvuxdsp | VSX Vector Convert Unsigned Dword to SP |
| 111100 .... 11111 .... $011101000 .$. | XX2 | I | 699 | v2.06 |  |  | xvcvuxwdp | VSX Vector Convert Unsigned Word to DP |
| 111100 $\ldots . .11111 \ldots . . .1010101000 .$. | XX2 | I | 699 | v2.06 |  |  | xvcvuxwsp | VSX Vector Convert Unsigned Word to SP |
| 111100 $\ldots \ldots . \ldots . . . . . . .01111000 \ldots$ | XX3 | 1 | 700 | v2.06 |  |  | xvdivdp | VSX Vector Divide DP |
| 111100 $\ldots . . . . . . . . . . . .01011000 \ldots$ | XX3 | I | 702 | v2.06 |  |  | xvdivsp | VSX Vector Divide SP |
| 111100 .... ......... $01100001 \ldots$ | XX3 | 1 | 705 | v2.06 |  |  | xvmaddadp | VSX Vector Multiply-Add Type-A DP |
| 111100 $\ldots \ldots . \ldots . . . . . . .01000001 \ldots$ | XX3 | I | 708 | v2.06 |  |  | xvmaddasp | VSX Vector Multiply-Add Type-A SP |
| 111100 $\ldots . . . . . . . . . . .01101001 \ldots$ | XX3 | I | 705 | v2.06 |  |  | xvmaddmdp | VSX Vector Multiply-Add Type-M DP |
| 111100 $\ldots \ldots . \ldots \ldots . . . . . .10101001 \ldots$ | XX3 | I | 708 | v2.06 |  |  | xvmaddmsp | VSX Vector Multiply-Add Type-M SP |
| 111100 $\ldots \ldots . \ldots \ldots \ldots 11100000 \ldots$ | XX3 | 1 | 711 | v2.06 |  |  | xvmaxdp | VSX Vector Maximum DP |
| 111100 $\ldots \ldots . \ldots . . . . . .11000000 \ldots$ | XX3 | 1 | 713 | v2.06 |  |  | xvmaxsp | VSX Vector Maximum SP |
| 111100 $\ldots \ldots . \ldots . . . . . . .11101000 \ldots$ | XX3 | I | 715 | v2.06 |  |  | xvmindp | VSX Vector Minimum DP |
| 111100 $\ldots \ldots \ldots \ldots \ldots 11001000 \ldots$ | XX3 | I | 717 | v2.06 |  |  | xvminsp | VSX Vector Minimum SP |
| 111100 $\ldots \ldots . . . . . . . . .01110001 \ldots$ | XX3 | I | 719 | v2.06 |  |  | xvmsubadp | VSX Vector Multiply-Subtract Type-A DP |
| 111100 $\ldots \ldots . \ldots . . . . . . .01010001 \ldots$ | XX3 | 1 | 722 | v2.06 |  |  | xvmsubasp | VSX Vector Multiply-Subtract Type-A SP |
| 111100 $\ldots . . . . . . . . . . .01111001 \ldots$ | XX3 | I | 719 | v2.06 |  |  | xvmsubmdp | VSX Vector Multiply-Subtract Type-M DP |
| 111100 $\ldots \ldots . \ldots . . . . . .01011001 \ldots$ | XX3 | 1 | 722 | v2.06 |  |  | xvmsubmsp | VSX Vector Multiply-Subtract Type-M SP |
| 111100 $\ldots \ldots . \ldots . . . . . . .01110000 \ldots$ | XX3 | 1 | 725 | v2.06 |  |  | xvmuldp | VSX Vector Multiply DP |
| 111100 $\ldots \ldots . \ldots . . . . . . .101010000 \ldots$ | XX3 | I | 727 | v2.06 |  |  | xvmulsp | VSX Vector Multiply SP |
| 111100 $\ldots . . .11111 \ldots . .111101001$. | XX2 | 1 | 729 | v2.06 |  |  | xvnabsdp | VSX Vector Negative Absolute DP |
| 111100 $\ldots . .11111 \ldots . . .110101001$. | XX2 | 1 | 729 | v2.06 |  |  | xvnabssp | VSX Vector Negative Absolute SP |
| 111100 $\ldots . . .11111 \ldots \ldots 111111001 .$. | XX2 | I | 730 | v2.06 |  |  | xvnegdp | VSX Vector Negate DP |
| 111100 .... 11111 .... 11011 1001.. | XX2 | 1 | 730 | v2.06 |  |  | xvnegsp | VSX Vector Negate SP |
| 111100 $\ldots . . . . . . . . . . .11100001 \ldots$ | XX3 | I | 731 | v2.06 |  |  | xvnmaddadp | VSX Vector Negative Multiply-Add Type-A DP |
| 1111100 $\ldots \ldots . \ldots . . . . . . .11000001 \ldots$ | XX3 | I | 736 | v2.06 |  |  | xvnmaddasp | VSX Vector Negative Multiply-Add Type-A SP |
| 111100 $\ldots \ldots . . . . . . . . . .111011001 \ldots$ | XX3 | 1 | 731 | v2.06 |  |  | xvnmaddmdp | VSX Vector Negative Multiply-Add Type-M DP |
| 111100 $\ldots \ldots \ldots . . . . . . .11001001 \ldots$ | XX3 | 1 | 736 | v2.06 |  |  | xvnmaddmsp | VSX Vector Negative Multiply-Add Type-M SP |
| 111100 $\ldots \ldots . . . . . . . . . .11110001 \ldots$ | XX3 | I | 739 | v2.06 |  |  | xvnmsubadp | VSX Vector Negative Multiply-Subtract Type-A DP |
| 111100 $\ldots \ldots . . . . . . . . .110101001 \ldots$ | XX3 | 1 | 742 | v2.06 |  |  | xvnmsubasp | VSX Vector Negative Multiply-Subtract Type-A SP |
| $111100 \ldots . . . . . . . . . . .11111001 \ldots$ | XX3 | I | 739 | v2.06 |  |  | xvnmsubmdp | VSX Vector Negative Multiply-Subtract Type-M DP |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 8 of 17)

| Instruction $^{\mathbf{1}}$11 11111 $11112 \quad 22222222233$ <br> $012345 \quad 67890$ 12345 $67890 \quad 12345 \quad 678901$ |  | $\begin{aligned} & \text { 두 } \\ & \text { ob } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \stackrel{0}{0} \\ & \text { N } \\ & \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 111100 $\ldots \ldots . . . . . . . . . .11011001 \ldots$ | XX3 | \| | 742 | v2.06 |  |  | xvnmsubmsp | VSX Vector Negative Multiply-Subtract Type-M SP |
| 111100 .... $11111 \ldots . . .011001001$. | XX2 | 1 | 745 | v2.06 |  |  | xvrdpi | VSX Vector Round DP to Integral to Nearest Away |
| 111100 $\ldots . .11111 \ldots \ldots 011101011$. | XX2 | 1 | 745 | v2.06 |  |  | xvrdpic | VSX Vector Round DP to Integral using Current rounding mode |
| $111100 \ldots . .11111 \ldots . .011111001$. | XX2 | I | 746 | v2.06 |  |  | xvrdpim | VSX Vector Round DP to Integral toward -Infinity |
| 111100 $\ldots . .11111 \ldots \ldots 011101001$. | XX2 | I | 746 | v2.06 |  |  | xvrdpip | VSX Vector Round DP to Integral toward +Infinity |
| 111100 $\ldots . .11111 \ldots \ldots 011011001$. | XX2 | I | 747 | v2.06 |  |  | xvrdpiz | VSX Vector Round DP to Integral toward Zero |
| $111100 \ldots . .11111 \ldots . .011011010$. | XX2 | I | 748 | v2.06 |  |  | xvredp | VSX Vector Reciprocal Estimate DP |
| 111100 $\ldots . .111111 \ldots . .1010011010$. | XX2 | I | 749 | v2.06 |  |  | xvresp | VSX Vector Reciprocal Estimate SP |
| 111100 $\ldots . .11111 \ldots \ldots 010001001$. | XX2 | I | 750 | v2.06 |  |  | xvrspi | VSX Vector Round SP to Integral to Nearest Away |
| 111100 $\ldots . .11111 \ldots . .010101011$. | XX2 | 1 | 750 | v2.06 |  |  | xvrspic | VSX Vector Round SP to Integral using Current rounding mode |
| 111100 $\ldots . .11111 \ldots . .010111001$. | XX2 | 1 | 751 | v2.06 |  |  | xvrspim | VSX Vector Round SP to Integral toward -Infinity |
| 111100 .... $11111 \ldots . . .010101001$. | XX2 | - | 751 | v2.06 |  |  | xvrspip | VSX Vector Round SP to Integral toward +Infinity |
| 111100 $\ldots . .11111 \ldots \ldots 010011001$. | XX2 | I | 752 | v2.06 |  |  | xvrspiz | VSX Vector Round SP to Integral toward Zero |
| 111100 $\ldots . .11111 \ldots \ldots 011001010$. | XX2 | I | 752 | v2.06 |  |  | xvrsqriedp | VSX Vector Reciprocal Square Root Estimate DP |
| 111100 $\ldots . .11111 \ldots \ldots 010001010$ | XX2 | 1 | 754 | v2.06 |  |  | xvrsqriesp | VSX Vector Reciprocal Square Root Estimate SP |
| 111100 $\ldots . .11111 \ldots . .011001011$. | XX2 | I | 755 | v2.06 |  |  | xvsqrtdp | VSX Vector Square Root DP |
| 111100 $\ldots . .11111 \ldots \ldots 010001011$. | XX2 | I | 756 | v2.06 |  |  | xvsqrtsp | VSX Vector Square Root SP |
| $111100 \ldots \ldots . . . . . . . . . .01101000 \ldots$ | XX3 | 1 | 757 | v2.06 |  |  | xvsubdp | VSX Vector Subtract DP |
| 111100 $\ldots \ldots \ldots \ldots \ldots .101001000 \ldots$ | XX3 | I | 759 | v2.06 |  |  | xvsubsp | VSX Vector Subtract SP |
| $111100 \ldots 11 \ldots . . . . . . .10111101 . .1$ | XX3 | I | 761 | v2.06 |  |  | xvtdivdp | VSX Vector Test for software Divide DP |
| $111100 \ldots 11 \ldots . . . . . . .1011101 .1$ | XX3 | I | 762 | v2.06 |  |  | xvtdivsp | VSX Vector Test for software Divide SP |
| 111100 $\ldots 1111111 \ldots \ldots 011101010.1$ | XX2 | 1 | 763 | v2.06 |  |  | xvtsqrtdp | VSX Vector Test for software Square Root DP |
| $111100 \ldots 1111111 \ldots . .010101010 .1$ | XX2 | I | 763 | v2.06 |  |  | xvtsqrisp | VSX Vector Test for software Square Root SP |
| 111100 $\ldots \ldots \ldots \ldots \ldots 10000$ 010........ | XX3 | I | 771 | v2.06 |  |  | xxland | VSX Vector Logical AND |
| 111100 $\ldots \ldots \ldots \ldots \ldots .10001010 \ldots$ | XX3 | I | 771 | v2.06 |  |  | xxlandc | VSX Vector Logical AND with Complement |
| 111100 $\ldots \ldots . \ldots \ldots \ldots . . . .10100 \quad 010 \ldots$ | XX3 | I | 773 | v2.06 |  |  | xxInor | VSX Vector Logical NOR |
| 111100 $\ldots \ldots . . . . . . . . . .10010010 \ldots$ | XX3 | I | 774 | v2.06 |  |  | xxlor | VSX Vector Logical OR |
| 111100 .............. 10011 010... | XX3 | 1 | 774 | v2.06 |  |  | xxIxor | VSX Vector Logical XOR |
| 111100 $\ldots \ldots . \ldots . . \ldots . .100010010 \ldots$ | XX3 | - | 775 | v2.06 |  |  | xxmrghw | VSX Vector Merge Word High |
|  | XX3 | 1 | 775 | v2.06 |  |  | xxmrglw | VSX Vector Merge Word Low |
| 111100 $\ldots \ldots . . . . . . . . . . . ~ 0 . . ~ 01010 . ~$ | XX3 | I | 777 | v2.06 |  |  | xxpermdi | VSX Vector Dword Permute Immediate |
|  | XX4 | I | 777 | v2.06 |  |  | xxsel | VSX Vector Select |
|  | XX3 | I | 778 | v2.06 |  |  | xxsldwi | VSX Vector Shift Left Double by Word Immediate |
| 111100 $\ldots . .111 . .1 . . .010100100$. | XX2 | I | 778 | v2.06 |  |  | xxspltw | VSX Vector Splat Word |
|  | X | I | 96 | v2.05 |  |  | cmpb | Compare Bytes |
| 111011............. 0000000010. | X | I | 195 | v2.05 |  |  | dadd[.] | DFP Add |
| 111111............. 0000000010. | X | I | 195 | v2.05 |  |  | daddq[.] | DFP Add Quad |
| 111111.... $11111 \ldots \ldots 1100100010$. | X | 1 | 217 | v2.05 |  |  | dcffixq[.] | DFP Convert From Fixed Quad |
| 111011...11 $\ldots . . \ldots \ldots 00100000101$ | X | I | 201 | v2.05 |  |  | dcmpo | DFP Compare Ordered |
| 111111...11 $\ldots . . . \ldots . .00100000101$ | X | I | 201 | v2.05 |  |  | dcmpoq | DFP Compare Ordered Quad |
| 111011...11 .......... 10100000101 | X | I | 200 | v2.05 |  |  | dcmpu | DFP Compare Unordered |
| 111111...11 $\ldots . . .1 . . .10100000101$ | X | I | 200 | v2.05 |  |  | dcmpuq | DFP Compare Unordered Quad |
| 111011.... $11111 \ldots . . .0100000010$. | X | I | 215 | V2.05 |  |  | dctdp[.] | DFP Convert To DFP Long |
| 111011..... $11111 \ldots \ldots .10100100010$. | X | 1 | 217 | v2.05 |  |  | dctix[.] | DFP Convert To Fixed |
| 111111.... $11111 \ldots . . .0100100010$. | X | I | 217 | v2.05 |  |  | dctfixq[.] | DFP Convert To Fixed Quad |
| 111111.... $11111 \ldots \ldots 0100000010$. | X | I | 215 | v2.05 |  |  | dctqpq[] | DFP Convert To DFP Extended |
| 111011......111 .... 0101000010. | X | 1 | 219 | v2.05 |  |  | ddedpd[.] | DFP Decode DPD To BCD |
| 111111......111.... 0101000010. | X | I | 219 | v2.05 |  |  | ddedpdq[.] | DFP Decode DPD To BCD Quad |
| 111011............. 1000100010. | X | I | 198 | v2.05 |  |  | ddiv[.] | DFP Divide |
| 111111............. 1000100010. | X | I | 198 | v2.05 |  |  | ddivq[.] | DFP Divide Quad |
| 111011..... $11111 \ldots . .1101000010$. | X | 1 | 219 | v2.05 |  |  | denbcd[.] | DFP Encode BCD To DPD |
| 111111.... .1111 .... 1101000010. | X | 1 | 219 | v2.05 |  |  | denbcdq[.] | DFP Encode BCD To DPD Quad |
| 111011 $\ldots \ldots . . . . . . . . . .1101100010$. | X | 1 | 220 | v2.05 |  |  | diex[.] | DFP Insert Exponent |
| 111111.............. 1101100010. | X | I | 220 | v2.05 |  |  | diexq[.] | DFP Insert Exponent Quad |
| 111011............. 0000100010. | X | I | 197 | v2.05 |  |  | dmul[.] | DFP Multiply |
| 111111............. 0000100010. | X | I | 197 | v2.05 |  |  | dmulq[.] | DFP Multiply Quad |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 9 of 17)


Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 10 of 17)

| Instruction $^{\mathbf{1}}$11 11111 11112 $22222 \quad 222233$ <br> $012345 \quad 67890$ 12345 67890 12345 <br> 678901    | 츨 | Y O 0 0 | $\begin{aligned} & \mathbb{0} \\ & \underset{\sim}{0} \end{aligned}$ |  |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | DS | I | 60 | v2.03 |  |  | stq | Store Qword |
| $011111 . . . . . . . . . . . . .00100001111$ | X | I | 247 | v2.03 |  |  | stvebx | Store Vector Element Byte Indexed |
| $011111 . \ldots . . . . . . . . . .00101001111$ | X | I | 247 | v2.03 |  |  | stvehx | Store Vector Element Hword Indexed |
| $011111 \ldots . . . . . . . . . . . . .00110001111$ | X | 1 | 248 | v2.03 |  |  | stvewx | Store Vector Element Word Indexed |
| $011111 \ldots . . . . . . . . . . . .00111001111$ | X | 1 | 248 | v2.03 |  |  | stvx | Store Vector Indexed |
|  | X | 1 | 248 | v2.03 |  |  | stvx\| | Store Vector Indexed Last |
| $0111111111111111 \ldots . .01000100101$ | X | 111 | 1038 | v2.03 | P | 64 | tlbiel | TLB Invalidate Entry Local |
|  | VX | I | 271 | v2.03 |  |  | vaddcuw | Vector Add \& Write Carry-Out Unsigned Word |
| $000100 \ldots \ldots . . . . . . . .00000001010$ | VX | 1 | 324 | v2.03 |  |  | vaddfp | Vector Add Floating-Point |
| $000100 \ldots \ldots . \ldots \ldots . .1100000000$ | VX | 1 | 271 | v2.03 |  |  | vaddsbs | Vector Add Signed Byte Saturate |
| $000100 \ldots \ldots . \ldots \ldots .01101000000$ | VX | I | 271 | v2.03 |  |  | vaddshs | Vector Add Signed Hword Saturate |
| $000100 \ldots \ldots . \ldots \ldots . .10 .110000000$ | VX | 1 | 272 | v2.03 |  |  | vaddsws | Vector Add Signed Word Saturate |
|  | VX | I | 272 | v2.03 |  |  | vaddubm | Vector Add Unsigned Byte Modulo |
| $000100 \ldots \ldots . \ldots \ldots . .101000000000$ | VX | I | 274 | v2.03 |  |  | vaddubs | Vector Add Unsigned Byte Saturate |
| $000100 \ldots \ldots . \ldots \ldots .00001000000$ | VX | 1 | 273 | v2.03 |  |  | vadduhm | Vector Add Unsigned Hword Modulo |
| $000100 \ldots \ldots . \ldots \ldots .01001000000$ | VX | 1 | 274 | v2.03 |  |  | vadduhs | Vector Add Unsigned Hword Saturate |
| $000100 \ldots \ldots . \ldots \ldots, \ldots 0010000000$ | VX | I | 273 | v2.03 |  |  | vadduwm | Vector Add Unsigned Word Modulo |
| $000100 \ldots \ldots . \ldots \ldots .01010000000$ | VX | I | 274 | v2.03 |  |  | vadduws | Vector Add Unsigned Word Saturate |
| $000100 \ldots \ldots . \ldots \ldots .10000000100$ | VX | I | 315 | v2.03 |  |  | vand | Vector Logical AND |
| $000100 \ldots \ldots . . . . . . . . .10001000100$ | VX | I | 315 | v2.03 |  |  | vandc | Vector Logical AND with Complement |
| $000100 \ldots \ldots . \ldots \ldots .10100000010$ | VX | I | 298 | v2.03 |  |  | vavgsb | Vector Average Signed Byte |
| $000100 \ldots \ldots . \ldots \ldots . . .10101000010$ | VX | I | 298 | v2.03 |  |  | vavgsh | Vector Average Signed Hword |
| $000100 \ldots \ldots . \ldots . . . . . . .10110000010$ | VX | I | 298 | v2.03 |  |  | vavgsw | Vector Average Signed Word |
| $000100 \ldots \ldots \ldots 10000000010$ | VX | I | 299 | v2.03 |  |  | vavgub | Vector Average Unsigned Byte |
| $000100 \ldots \ldots . \ldots \ldots .10001000010$ | VX | I | 299 | v2.03 |  |  | vavguh | Vector Average Unsigned Hword |
| $000100 \ldots \ldots . \ldots \ldots .10010000010$ | VX | I | 299 | v2.03 |  |  | vavguw | Vector Average Unsigned Word |
|  | VX | I | 328 | v2.03 |  |  | vcfsx | Vector Convert From Signed Word |
| $000100 \ldots \ldots . . . . . . . . . .01100001010$ | VX | I | 328 | v2.03 |  |  | vcfux | Vector Convert From Unsigned Word |
| $000100 \ldots \ldots . . . . . . . . . .1111000110$ | VC | I | 331 | v2.03 |  |  | vcmpbfp[.] | Vector Compare Bounds Floating-Point |
| $000100 \ldots \ldots . \ldots . . . . . . . . . . .0011000110$ | VC | I | 332 | v2.03 |  |  | vcmpeafp[.] | Vector Compare Equal To Floating-Point |
|  | VC | I | 306 | v2.03 |  |  | vcmpequb[.] | Vector Compare Equal Unsigned Byte |
|  | VC | I | 306 | v2.03 |  |  | vcmpequh[.] | Vector Compare Equal Unsigned Hword |
|  | VC | I | 307 | v2.03 |  |  | vcmpequw[.] | Vector Compare Equal Unsigned Word |
|  | VC | I | 332 | v2.03 |  |  | vcmpgefp[.] | Vector Compare Greater Than or Equal To Floating-Point |
| $000100 \ldots \ldots . . . . . . . . . .1011000110$ | VC | I | 333 | v2.03 |  |  | vcmpgtfp[.] | Vector Compare Greater Than Floating-Point |
| $000100 \ldots \ldots . \ldots . . . .1100000110$ | VC | I | 308 | v2.03 |  |  | vcmpgtsb[.] | Vector Compare Greater Than Signed Byte |
| $000100 \ldots \ldots . \ldots . . . . . . .1101000110$ | VC | I | 309 | v2.03 |  |  | vcmpgtsh[.] | Vector Compare Greater Than Signed Hword |
| $000100 \ldots \ldots . \ldots . . . . . . .1110000110$ | VC | I | 309 | v2.03 |  |  | vcmpgtsw[.] | Vector Compare Greater Than Signed Word |
| $000100 \ldots \ldots . \ldots . . . . . .1000000110$ | VC | I | 310 | v2.03 |  |  | vcmpgtub[.] | Vector Compare Greater Than Unsigned Byte |
| $000100 \ldots \ldots . \ldots . . . . . .1001000110$ | VC | I | 311 | v2.03 |  |  | vcmpgtuh[.] | Vector Compare Greater Than Unsigned Hword |
| $000100 \ldots \ldots . \ldots . . . . . . . .1010000110$ | VC | I | 311 | v2.03 |  |  | vcmpgtuw[.] | Vector Compare Greater Than Unsigned Word |
|  | VX | I | 327 | v2.03 |  |  | vctsxs | Vector Convert To Signed Word Saturate |
|  | VX | I | 327 | v2.03 |  |  | vctuxs | Vector Convert To Unsigned Word Saturate |
| $000100 \ldots . .11111 \ldots . .100110001010$ | VX | I | 334 | v2.03 |  |  | vexptefp | Vector 2 Raised to the Exponent Estimate Floating-Point |
| $000100 \ldots \ldots . \ldots . . . . . .00111001010$ | VX | I | 334 | v2.03 |  |  | vlogefp | Vector Log Base 2 Estimate Floating-Point |
| $000100 \ldots \ldots . . . . . . . . . . . . . . . .101110$ | VA | । | 325 | v2.03 |  |  | vmaddfp | Vector Multiply-Add Floating-Point |
| $000100 \ldots \ldots . \ldots . . . . .10000001010$ | VX | I | 326 | v2.03 |  |  | vmaxfp | Vector Maximum Floating-Point |
| $000100 \ldots \ldots . \ldots \ldots . . .100100000010$ | VX | I | 302 | v2.03 |  |  | vmaxsb | Vector Maximum Signed Byte |
|  | VX | I | 303 | v2.03 |  |  | vmaxsh | Vector Maximum Signed Hword |
| $000100 \ldots \ldots . \ldots \ldots .00110000010$ | VX | I | 303 | v2.03 |  |  | vmaxsw | Vector Maximum Signed Word |
| $000100 \ldots \ldots . \ldots \ldots .00000000010$ | VX | I | 302 | v2.03 |  |  | vmaxub | Vector Maximum Unsigned Byte |
| $000100 \ldots \ldots . \ldots . . . . .100001000010$ | VX | I | 303 | v2.03 |  |  | vmaxuh | Vector Maximum Unsigned Hword |
| $000100 \ldots \ldots . \ldots \ldots .00010000010$ | VX | 1 | 303 | v2.03 |  |  | vmaxuw | Vector Maximum Unsigned Word |
| $000100 \ldots \ldots . . . . . . . . . . . . . . . .100000$ | VA | 1 | 287 | v2.03 |  |  | vmhaddshs | Vector Multiply-High-Add Signed Hword Saturate |
| $000100 \ldots \ldots . . . . . . . . . . . . . . .100001$ | VA | I | 287 | v2.03 |  |  | vmhraddshs | Vector Multiply-High-Round-Add Signed Hword Saturate |
| $000100 \ldots \ldots . . . . . . . . .10001001010$ | VX | I | 326 | v2.03 |  |  | vminfp | Vector Minimum Floating-Point |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 11 of 17)

| Instruction ${ }^{1}$ |  |  |  |  | $\begin{aligned} & \text { ̌i } \\ & \text { on } \end{aligned}$ |  | $\begin{aligned} & \text { N } \\ & \stackrel{.0}{\omega} \\ & \frac{0}{\omega} \\ & \hline \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 012345 | 6789011234567890 | $\begin{aligned} & 22222 \\ & 12345 \end{aligned}$ | 622233 |  |  |  |  |  |  |  |  |
| 000100 | ............... | 01100 | 000010 | VX | 1 | 304 | v2.03 |  |  | vminsb | Vector Minimum Signed Byte |
| 000100 | ..... .......... | 01101 | 000010 | VX | I | 305 | v2.03 |  |  | vminsh | Vector Minimum Signed Hword |
| 000100 | $\ldots \ldots \ldots \ldots$ | 01110 | 000010 | VX | 1 | 305 | v2.03 |  |  | vminsw | Vector Minimum Signed Word |
| 000100 | $\ldots \ldots \ldots$ | 01000 | 000010 | VX | 1 | 304 | v2.03 |  |  | vminub | Vector Minimum Unsigned Byte |
| 000100 | .............. | 01001 | 000010 | VX | I | 305 | v2.03 |  |  | vminuh | Vector Minimum Unsigned Hword |
| 000100 |  | 01010 | 000010 | VX | I | 305 | v2.03 |  |  | vminuw | Vector Minimum Unsigned Word |
| 000100 | .... ......... | .... | 100010 | VA | 1 | 288 | v2.03 |  |  | vmladduhm | Vector Multiply-Low-Add Unsigned Hword Modulo |
| 000100 | .... .......... | 00000 | 001100 | VX | I | 257 | v2.03 |  |  | vmrghb | Vector Merge High Byte |
| 000100 | $\ldots \ldots \ldots$ | 00001 | 001100 | VX | I | 257 | v2.03 |  |  | vmrghh | Vector Merge High Hword |
| 000100 | .............. | 00010 | 001100 | VX | 1 | 258 | v2.03 |  |  | vmrghw | Vector Merge High Word |
| 000100 | .... .......... | 00100 | 001100 | VX | I | 257 | v2.03 |  |  | vmrglb | Vector Merge Low Byte |
| 000100 | ............. | 00101 | 001100 | VX | I | 257 | v2.03 |  |  | vmrglh | Vector Merge Low Hword |
| 000100 | ..... ........... | 00110 | 001100 | VX | I | 258 | v2.03 |  |  | vmrglw | Vector Merge Low Word |
| 000100 | $\ldots \ldots \ldots \ldots$ | ..... | 100101 | VA | 1 | 289 | v2.03 |  |  | vmsummbm | Vector Multiply-Sum Mixed Byte Modulo |
| 000100 | ..... ........... | ..... | 101000 | VA | 1 | 289 | v2.03 |  |  | vmsumshm | Vector Multiply-Sum Signed Hword Modulo |
| 000100 | ..... .......... | ..... | 101001 | VA | I | 290 | v2.03 |  |  | vmsumshs | Vector Multiply-Sum Signed Hword Saturate |
| 000100 | $\ldots . . . . . . .$. | ..... | 100100 | VA | I | 288 | v2.03 |  |  | vmsumubm | Vector Multiply-Sum Unsigned Byte Modulo |
| 000100 | ..... .......... | ..... | 100110 | VA | 1 | 290 | v2.03 |  |  | vmsumuhm | Vector Multiply-Sum Unsigned Hword Modulo |
| 000100 | ............... | ..... | 100111 | VA | 1 | 291 | v2.03 |  |  | vmsumuhs | Vector Multiply-Sum Unsigned Hword Saturate |
| 000100 | $\ldots \ldots \ldots \ldots$ | 01100 | 001000 | VX | I | 283 | v2.03 |  |  | vmulesb | Vector Multiply Even Signed Byte |
| 000100 | ................ | 01101 | 001000 | VX | 1 | 284 | v2.03 |  |  | vmulesh | Vector Multiply Even Signed Hword |
| 000100 | ..... .......... | 01000 | 001000 | VX | 1 | 283 | v2.03 |  |  | vmuleub | Vector Multiply Even Unsigned Byte |
| 000100 | $\ldots \ldots$ | 01001 | 001000 | VX | 1 | 284 | v2.03 |  |  | vmuleuh | Vector Multiply Even Unsigned Hword |
| 000100 | .... .......... | 00100 | 001000 | VX | I | 283 | v2.03 |  |  | vmulosb | Vector Multiply Odd Signed Byte |
| 000100 | ..... .......... | 00101 | 001000 | VX | 1 | 284 | v2.03 |  |  | vmulosh | Vector Multiply Odd Signed Hword |
| 000100 | , | 00000 | 001000 | VX | 1 | 283 | v2.03 |  |  | vmuloub | Vector Multiply Odd Unsigned Byte |
| 000100 | ..... ........... | 00001 | 001000 | VX | 1 | 284 | v2.03 |  |  | vmulouh | Vector Multiply Odd Unsigned Hword |
| 000100 | ............. |  | 101111 | VA | 1 | 325 | v2.03 |  |  | vnmsubfp | Vector Negative Multiply-Subtract Floating-Point |
| 000100 | ........... | 10100 | 000100 | VX | 1 | 316 | v2.03 |  |  | vnor | Vector Logical NOR |
| 000100 | $\ldots \ldots \ldots \ldots$ | 10010 | 000100 | VX | 1 | 316 | v2.03 |  |  | vor | Vector Logical OR |
| 000100 | ..... .......... | ..... | 101011 | VA | 1 | 262 | v2.03 |  |  | vperm | Vector Permute |
| 000100 | ............ | 01100 | 001110 | VX | 1 | 250 | v2.03 |  |  | vpkpx | Vector Pack Pixel |
| 000100 | ..... .......... | 00110 | 001110 | VX | 1 | 251 | v2.03 |  |  | vpkshss | Vector Pack Signed Hword Signed Saturate |
| 000100 | ..... .......... | 00100 | 001110 | VX | I | 252 | v2.03 |  |  | vpkshus | Vector Pack Signed Hword Unsigned Saturate |
| 000100 | ........ ... | 00111 | 001110 | VX | 1 | 252 | v2.03 |  |  | vpkswss | Vector Pack Signed Word Signed Saturate |
| 000100 | ............. | 00101 | 001110 | VX | 1 | 253 | v2.03 |  |  | vpkswus | Vector Pack Signed Word Unsigned Saturate |
| 000100 | $\ldots . . . . . . .$. | 00000 | 001110 | VX | I | 253 | v2.03 |  |  | vpkuhum | Vector Pack Unsigned Hword Unsigned Modulo |
| 000100 | ..... ........ | 00010 | 001110 | VX | 1 | 254 | v2.03 |  |  | vpkuhus | Vector Pack Unsigned Hword Unsigned Saturate |
| 000100 | ............. | 00001 | 001110 | VX | 1 | 254 | v2.03 |  |  | vpkuwum | Vector Pack Unsigned Word Unsigned Modulo |
| 000100 | $\ldots$ | 00011 | 001110 | VX | I | 254 | v2.03 |  |  | vpkuwus | Vector Pack Unsigned Word Unsigned Saturate |
| 000100 | ..... \|1111 ..... | 00100 | 001010 | VX | 1 | 335 | v2.03 |  |  | vrefp | Vector Reciprocal Estimate Floating-Point |
| 000100 | ..... 11111..... | 01011 | 001010 | VX | 1 | 329 | v2.03 |  |  | vrfim | Vector Round to Floating-Point Integral toward -Infinity |
| 000100 | .....11111.... | 01000 | 001010 | VX | 1 | 329 | v2.03 |  |  | vrfin | Vector Round to Floating-Point Integral Nearest |
| 000100 | ..... 11111 ..... | 01010 | 001010 | VX | I | 329 | v2.03 |  |  | vrip | Vector Round to Floating-Point Integral toward +Infinity |
| 000100 | ..... 11111..... | 01001 | 001010 | VX | 1 | 330 | v2.03 |  |  | vrif | Vector Round to Floating-Point Integral toward Zero |
| 000100 | ... .......... | 00000 | 000100 | VX | 1 | 318 | v2.03 |  |  | vrlb | Vector Rotate Left Byte |
| 000100 | $\ldots \ldots \ldots \ldots$ | 00001 | 000100 | VX | 1 | 318 | v2.03 |  |  | vrlh | Vector Rotate Left Hword |
| 000100 | ............... | 00010 | 000100 | VX | I | 318 | v2.03 |  |  | vrlw | Vector Rotate Left Word |
| 000100 | ..... \|l||1..... | 00101 | 001010 | VX | 1 | 335 | v2.03 |  |  | visqritefp | Vector Reciprocal Square Root Estimate Floating-Point |
| 000100 | .... .......... | ... | 101010 | VA | 1 | 263 | v2.03 |  |  | vsel | Vector Select |
| 000100 | .... .......... | 00111 | 000100 | VX | I | 266 | v2.03 |  |  | vsl | Vector Shift Left |
| 000100 | .... .......... | 00100 | 000100 | VX | 1 | 319 | v2.03 |  |  | vslb | Vector Shit Left Byte |
| 000100 | ............. | 1.... | 101100 | VA | 1 | 265 | v2.03 |  |  | vsldoi | Vector Shitt Left Double by Octet Immediate |
| 000100 | $\ldots$ | 00101 | 000100 | VX | I | 319 | v2.03 |  |  | vslh | Vector Shift Left Hword |
| 000100 | ..... .......... | 10000 | 001100 | VX | 1 | 266 | v2.03 |  |  | vslo | Vector Shitt Left by Octet |
| 000100 | ..... .......... | 00110 | 000100 | VX | I | 319 | v2.03 |  |  | vslw | Vector Shitt Left Word |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 12 of 17)

| Instruction $^{\mathbf{1}}$ $012345 \quad 678901111111112 \quad 22222 \quad 222233$ $01234567890 \quad 12345678901$ |  | $\begin{aligned} & \text { Y } \\ & \text { o } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & \underset{\sim}{0} \end{aligned}$ |  |  | $\begin{aligned} & \sigma_{0} \\ & 0 \\ & 0 \\ & 0 \\ & 0 \\ & 0 \\ & \end{aligned}$ | 0 <br> 0 <br> 0 <br> 1 <br> 1 | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $000100 \ldots . .1 \ldots \ldots .101000001100$ | VX | I | 260 | v2.03 |  |  | vspltb | Vector Splat Byte |
| $000100 \ldots . .11 \ldots \ldots .101001001100$ | VX | I | 260 | v2.03 |  |  | vsplth | Vector Splat Hword |
| $000100 \ldots \ldots . . .1111101100001100$ | VX | I | 261 | v2.03 |  |  | vspltisb | Vector Splat Immediate Signed Byte |
| $000100 \ldots \ldots . . .11111101101001100$ | VX | 1 | 261 | v2.03 |  |  | vspltish | Vector Splat Immediate Signed Hword |
| $000100 \ldots \ldots \ldots 1111101110001100$ | VX | I | 261 | v2.03 |  |  | vspltisw | Vector Splat Immediate Signed Word |
| $000100 \ldots \ldots 111 . . \ldots . .01010001100$ | VX | I | 260 | v2.03 |  |  | vspltw | Vector Splat Word |
| $000100 \ldots \ldots . . . . . . . .01011000100$ | VX | I | 266 | v2.03 |  |  | vsr | Vector Shift Right |
|  | VX | I | 321 | v2.03 |  |  | vsrab | Vector Shift Right Algebraic Byte |
| $000100 \ldots \ldots . . . . . . . . .01101000100$ | VX | I | 321 | v2.03 |  |  | vsrah | Vector Shift Right Algebraic Hword |
|  | VX | 1 | 321 | v2.03 |  |  | vsraw | Vector Shift Right Algebraic Word |
| $000100 \ldots \ldots . \ldots \ldots .01000000100$ | VX | I | 320 | v2.03 |  |  | vsrb | Vector Shift Right Byte |
| $000100 \ldots \ldots . \ldots \ldots . .101001000100$ | VX | 1 | 320 | v2.03 |  |  | vsrh | Vector Shift Right Hword |
| $000100 \ldots \ldots . . . . . . . .10001001100$ | VX | I | 266 | v2.03 |  |  | vsro | Vector Shift Right by Octet |
|  | VX | I | 320 | v2.03 |  |  | VSTw | Vector Shift Right Word |
| $000100 \ldots \ldots . \ldots \ldots .10110000000$ | VX | I | 277 | v2.03 |  |  | vsubcuw | Vector Subtract \& Write Carry-Out Unsigned Word |
|  | VX | I | 324 | v2.03 |  |  | vsubfp | Vector Subtract Floating-Point |
| $000100 \ldots \ldots \ldots \ldots .11100000000$ | VX | I | 277 | v2.03 |  |  | vsubsbs | Vector Subtract Signed Byte Saturate |
| $000100 \ldots \ldots . \ldots \ldots .11101000000$ | VX | I | 277 | v2.03 |  |  | vsubshs | Vector Subtract Signed Hword Saturate |
| $000100 \ldots \ldots \ldots \ldots .11110000000$ | VX | I | 278 | v2.03 |  |  | vsubsws | Vector Subtract Signed Word Saturate |
| $000100 \ldots \ldots . . . . . . . .10000000000$ | VX | I | 279 | v2.03 |  |  | vsububm | Vector Subtract Unsigned Byte Modulo |
| $000100 \ldots \ldots . \ldots \ldots .11000000000$ | VX | I | 280 | v2.03 |  |  | vsububs | Vector Subtract Unsigned Byte Saturate |
| $000100 \ldots \ldots \ldots \ldots .10001000000$ | VX | I | 279 | v2.03 |  |  | vsubuhm | Vector Subtract Unsigned Hword Modulo |
| $000100 \ldots \ldots . \ldots \ldots .11001000000$ | VX | I | 280 | v2.03 |  |  | vsubuhs | Vector Subtract Unsigned Hword Saturate |
| $000100 \ldots \ldots \ldots 10010000000$ | VX | I | 279 | v2.03 |  |  | vsubuwm | Vector Subtract Unsigned Word Modulo |
| $000100 \ldots \ldots \ldots \ldots .11010000000$ | VX | I | 280 | v2.03 |  |  | vsubuws | Vector Subtract Unsigned Word Saturate |
| $000100 \ldots \ldots . \ldots \ldots .11010001000$ | VX | 1 | 292 | v2.03 |  |  | vsum2sws | Vector Sum across Half Signed Word Saturate |
| $000100 \ldots \ldots . . . . . . . . .11100001000$ | VX | I | 293 | v2.03 |  |  | vsum4sbs | Vector Sum across Quarter Signed Byte Saturate |
| $000100 \ldots \ldots . . . . . . . . .11001001000$ | VX | I | 293 | v2.03 |  |  | vsum4shs | Vector Sum across Quarter Signed Hword Saturate |
| $000100 \ldots \ldots . \ldots . . . . .11000001000$ | VX | I | 294 | v2.03 |  |  | vsum4ubs | Vector Sum across Quarter Unsigned Byte Saturate |
| $000100 \ldots \ldots . . . . . . . . . .11110001000$ | VX | 1 | 292 | v2.03 |  |  | vsumsws | Vector Sum across Signed Word Saturate |
| $000100 \ldots . .11111 \ldots . .01101001110$ | VX | I | 255 | v2.03 |  |  | vupkhpx | Vector Unpack High Pixel |
| $000100 \ldots . .11111 \ldots . .01000001110$ | VX | I | 256 | v2.03 |  |  | vupkhsb | Vector Unpack High Signed Byte |
| $000100 \ldots . .11111 \ldots . .01001001110$ | VX | I | 256 | v2.03 |  |  | vupkhsh | Vector Unpack High Signed Hword |
| $000100 \ldots . .11111 \ldots . .01111001110$ | VX | I | 255 | v2.03 |  |  | vupklpx | Vector Unpack Low Pixel |
| $000100 \ldots . .11111 \ldots . .01010001110$ | VX | I | 256 | v2.03 |  |  | vupklsb | Vector Unpack Low Signed Byte |
| $000100 \ldots . .11111 \ldots . .01011001110$ | VX | I | 256 | v2.03 |  |  | vupklsh | Vector Unpack Low Signed Hword |
| $000100 \ldots \ldots . . . . . . . . .10011000100$ | VX | I | 316 | v2.03 |  |  | vxor | Vector Logical XOR |
| 111111..... $11111 \ldots . .1111111000$. | A | 1 | 155 | v2.02 |  |  | fre[.] | Floating Reciprocal Estimate |
| 111111.... $11111 \ldots . .0111101000$. | X | I | 167 | v2.02 |  |  | frim[.] | Floating Round To Integer Minus |
| 111111.... $11111 \ldots . .0110001000$. | X | I | 167 | v2.02 |  |  | frin[.] | Floating Round To Integer Nearest |
| 111111.... $11111 \ldots . . .0111001000$. | X | I | 167 | v2.02 |  |  | frip[.] | Floating Round To Integer Plus |
| 111111.... $11111 \ldots . .0110101000$. | X | I | 167 | v2.02 |  |  | friz[.] | Floating Round To Integer Zero |
| 111011 $\ldots . .11111 \ldots . .1111111010$. | A | 1 | 156 | v2.02 |  |  | frsqrtes[.] | Floating Reciprocal Square Root Estimate Single |
| 010011 \|1111 11111 11111 01000100101 | XL | 111 | 955 | v2.02 | H |  | hrfid | Return From Interrupt Dword Hypervisor |
| 011111......... 11111 00011110101 | X | I | 96 | v2.02 |  |  | popentb | Population Count Byte |
| 011111..... 1...... 100000100111 | XFX | I | 121 | v2.01 |  |  | mfocrf | Move From One CR Field |
| $011111 \ldots . .1 \ldots \ldots 100100100001$ | XFX | I | 120 | v2.01 |  |  | mtocrf | Move To One CR Field |
| 011111..... $11111 \ldots \ldots .11100100111$ | X | 111 | 1030 | v2.00 | P |  | slbmfee | SLB Move From Entry ESID |
| 011111..... $11111 \ldots . . .11010100111$ | X | 111 | 1030 | v2.00 | P |  | slbmfev | SLB Move From Entry VSID |
| $011111 \ldots . .11111 \ldots . .01100100101$ | X | 111 | 1029 | v2.00 | P |  | slbmte | SLB Move To Entry |
|  | DQ | 1 | 59 | v2.03 |  |  | Iq | Load Qword |
| 01001111111111111111100010100101 | XL | 111 | 953 | v3.0 | P |  | rfscV | Return From System Call Vectored |
| $01000111111111111111 . \ldots . .11101$ | SC | 1 | 43 | v3.0 |  |  | SCV | System Call Vectored |
| 011111......... 111110000111010. | X | 1 | 98 | PPC |  | SR | cntlzd[.] | Count Leading Zeros Dword |
| $011111111 \ldots \ldots . . . . . . . .00010101101$ | X | 11 | 854 | PPC |  |  | dcbf | Data Cache Block Flush |
| $01111111111 \ldots . . . . . . . .00001101101$ | X | 11 | 853 | PPC |  |  | dcbst | Data Cache Block Store |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 13 of 17)

| Instruction ${ }^{\mathbf{1}}$ <br> 111111111222222222233 <br> 0123456789012345678912345678901 |  |  | $\begin{aligned} & \mathbb{\otimes} \\ & \stackrel{\pi}{2} \end{aligned}$ | $\begin{aligned} & N \\ & \text { N } \\ & \text { N } \\ & \text { Non } \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 011111 .......... ..... 01000101101 | X | 11 | 851 | PPC |  |  | dcbt | Data Cache Block Touch |
| 011111 ..... ......... 00111101101 | X | 11 | 852 | PPC |  |  | dcbist | Data Cache Block Touch for Store |
|  | XO | 1 | 82 | PPC |  | SR | divd[0][.] | Divide Dword |
|  | XO | 1 | 82 | PPC |  | SR | divdu[0][] | Divide Dword Unsigned |
|  | XO | 1 | 75 | PPC |  | SR | diww[0][] | Divide Word |
| $011111 . . . .{ }^{\text {a }}$.... ..... . 1111001011. | XO | 1 | 75 | PPC |  | SR | diwul[0][] | Divide Word Unsigned |
| 011111 \||||| ||1|| ||||| 11010101101 | X | 11 | 879 | PPC |  |  | eieio | Enforce In-order Execution of I/O |
|  | X | 1 | 94 | PPC |  | SR | extsb[.] | Extend Sign Byte |
| 011111.......... 111111111011010. | X | I | 98 | PPC |  | SR | extsw[] | Extend Sign Word |
| 111011............... \||||| 10101. | A | 1 | 153 | PPC |  |  | fadds[.] | Floating Add Single |
|  | X | 1 | 164 | PPC |  |  | fcfid[.] | Floating Convert From Integer Dword |
|  | X | I | 160 | PPC |  |  | fctid[] | Floating Convert To Integer Dword |
|  | X | I | 161 | PPC |  |  | ftidid[.] | Floating Convert To Integer Dword truncate |
| 111011.............. \|111| 10010. | A | 1 | 154 | PPC |  |  | fdivs[] | Floating Divide Single |
| 111011.................... 11101. | A | I | 158 | PPC |  |  | fmadds[.] | Floating Multiply-Add Single |
| 111011..... ..... ..... ..... 11100. | A | I | 159 | PPC |  |  | fmsubs[] | Floating Multiply-Subtract Single |
| 111011......... \|1111 ..... 11001. | A | 1 | 154 | PPC |  |  | fmuls[]. | Floating Multiply Single |
|  | A | I | 159 | PPC |  |  | fnmadds[] | Floating Negative Multiply-Add Single |
|  | A | 1 | 159 | PPC |  |  | fnmsubs[] | Floating Negative Multiply-Subtract Single |
| 111011.... \|1111 ..... |11|| 11000. | A | I | 155 | PPC |  |  | fres[] | Floating Reciprocal Estimate Single |
| 1111111..... \|1111 ..... |111| 11010. | A | 1 | 156 | PPC |  |  | frsqrte[.] | Floating Reciprocal Square Root Estimate |
|  | A | 1 | 169 | PPC |  |  | fsel[.] | Floating Select |
| 111011..... \||1|| ..... ||||| 10110. | A | 1 | 155 | PPC |  |  | fsqris[] | Floating Square Root Single |
| 111011..... .......... \||||| 10100. | A | 1 | 153 | PPC |  |  | fsubs[.] | Floating Subtract Single |
| 011111 \|1111 ......... 11110101101 | X | 11 | 842 | PPC |  |  | icbi | Instruction Cache Block Invalidate |
| 111010 ..... ..... ..... ..... .... 00 | DS | 1 | 53 | PPC |  |  | Id | Load Dword |
|  | X | 11 | 873 | PPC |  |  | Idarx | Load Dword And Reserve Indexed |
| 111010 ..... ..... ..... ..... .... 01 | DS | 1 | 53 | PPC |  |  | Idu | Load Dword with Update |
| $011111 \ldots . . . . . . . . . . .00001101011$ | X | 1 | 53 | PPC |  |  | Idux | Load Dword with Update Indexed |
|  | X | 1 | 53 | PPC |  |  | Idx | Load Dword Indexed |
| 111010 ..... ..... ..... ..... .... 10 | DS | I | 52 | PPC |  |  | Iwa | Load Word Algebraic |
| $011111 \ldots \ldots . . . . . . . .000000101001$ | X | 11 | 869 | PPC |  |  | Iwarx | Load Word \& Reserve Indexed |
|  | X | 1 | 52 | PPC |  |  | Iwaux | Load Word Algebraic with Update Indexed |
| 011111 ..... ..... ..... 01010101011 | X | I | 52 | PPC |  |  | Iwax | Load Word Algebraic Indexed |
| 011111 .............. 01011100111 | X | 11 | 902 | PPC |  |  | mftb | Move From Time Base |
| $011111 \ldots . . .1\|11.1\| 11100101100101$ | X | 111 | 978 | PPC | P |  | mtmsrd | Move To MSR Dword |
|  | XO | I | 80 | PPC |  | SR | mulhd[.] | Multiply High Dword |
| $011111 . . . . . . . . . . . . .1000001001$. | XO | I | 80 | PPC |  | SR | mulhdu[.] | Multiply High Dword Unsigned |
|  | XO | 1 | 74 | PPC |  | SR | mulhw[] | Multiply High Word |
| $011111 . . . .1 . . . . . . . .1000001011$. | XO | I | 74 | PPC |  | SR | mulhwu[] | Multiply High Word Unsigned |
| 011111 ..... ..... ..... . 011101001. | XO | 1 | 80 | PPC |  | SR | mulld[0][]] | Multiply Low Dword |
| 010011 \|11|| |1|1| |11|1 00000100101 | XL | 111 | 954 | PPC | P |  | rfid | Return from Interrupt Dword |
| 011110 ..... ..... ..... ..... 1000. | MDS | 1 | 103 | PPC |  | SR | rldcl[.] | Rotate Left Dword then Clear Left |
|  | MDS | 1 | 103 | PPC |  | SR | rldcr[.] | Rotate Left Dword then Clear Right |
| 011110 ..... ..... ..... ..... 010. | MD | I | 104 | PPC |  | SR | rldic[].] | Rotate Left Dword Immediate then Clear |
|  | MD | 1 | 104 | PPC |  | SR | rldicl[.] | Rotate Left Dword Immediate then Clear Left |
|  | MD | 1 | 105 | PPC |  | SR | rldicr[.] | Rotate Left Dword Immediate then Clear Right |
| 011110 ..... ..... ..... ..... . 011. | MD | 1 | 105 | PPC |  | SR | rldimi[.] | Rotate Left Dword Immediate then Mask Insert |
| $01000111111111111111 . . . . . .11 / 11$ | SC | I | 43 | PPC |  |  | sc | System Call |
| 011111 \|1... |11|| |1||| 01111100101 | X | 111 | 1027 | PPC | P |  | slbia | SLB Invalidate All |
| 011111 \|1111 |1111..... 01101100101 | X | 111 | 1024 | PPC | P |  | slbie | SLB Invalidate Entry |
|  | X | I | 108 | PPC |  | SR | sld[].] | Shift Left Dword |
| $011111 \ldots \ldots . . . . . . . . .1100011010$. | X | 1 | 109 | PPC |  | SR | srad[] | Shift Right Algebraic Dword |
| 011111............... 110011101. | XS | 1 | 109 | PPC |  | SR | sradi[.] | Shift Right Algebraic Dword Immediate |
|  | X | 1 | 108 | PPC |  | SR | srd[.] | Shift Right Dword |
| 111110 ............... ..... .... 00 | DS | 1 | 58 | PPC |  |  | std | Store Dword |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 14 of 17)

| Instruction $^{\mathbf{1}}$111111 11112 $22222 \quad 222233$  <br> $012345 \quad 67890$ 12345 67890 12345 <br> 678901    |  | Y O 0 0 | $\begin{aligned} & \mathbb{0} \\ & \underset{\sim}{0} \end{aligned}$ | $\begin{aligned} & N \\ & \text { N } \\ & \hline \frac{0}{0} \\ & \frac{0}{0} \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $011111 \ldots . . . . . . . . . . . .00110101101$ | X | 11 | 873 | PPC |  |  | stdcx. | Store Dword Conditional Indexed \& record |
| 111110 ..... .... ..... ..... .... 01 | DS | I | 58 | PPC |  |  | stdu | Store Dword with Update |
| 011111.............. 00101101011 | X | I | 58 | PPC |  |  | stdux | Store Dword with Update Indexed |
| $011111 \ldots . . . . . . . . . . . .00100101011$ | X | I | 58 | PPC |  |  | stdx | Store Dword Indexed |
|  | X | I | 148 | PPC |  |  | stfiwx | Store Floating as Integer Word Indexed |
| 011111.............. 00100101101 | X | 11 | 872 | PPC |  |  | stwcx. | Store Word Conditional Indexed \& record |
| 011111.............. 000101000. | XO | I | 70 | PPC |  | SR | subf[0][.] | Subtract From |
| 011111............. 00010001001 | X | I | 90 | PPC |  |  | td | Trap Dword |
| 000010 | D | I | 90 | PPC |  |  | tdi | Trap Dword Immediate |
| 011111 \||1|| 1|||| |11|| 10001 10110| | X | 111 | 1042 | PPC | H |  | tlbsync | TLB Synchronize |
| 111111.... $11111 \ldots . .0000001110$. | X | I | 162 | P2 |  |  | fctiw[.] | Floating Convert To Integer Word |
| 111111.... $11111 \ldots . .0000001111$. | X | I | 163 | P2 |  |  | fctiwz[.] | Floating Convert To Integer Word truncate |
| 111111.... $111111 \ldots . .1111110110$. | A | I | 155 | P2 |  |  | fsqri[.] | Floating Square Root |
| 011111.............. . 100001010. | XO | I | 70 | P1 |  | SR | add[0][.] | Add |
|  | XO | I | 71 | P1 |  | SR | addc[0][.] | Add Carrying |
| 011111.............. . 010001010. | XO | 1 | 72 | P1 |  | SR | adde[0][.] | Add Extended |
|  | D | I | 68 | P1 |  |  | addi | Add Immediate |
|  | D | I | 70 | P1 |  | SR | addic | Add Immediate Carrying |
|  | D | I | 70 | P1 |  | SR | addic. | Add Immediate Carrying \& record |
|  | D | I | 68 | P1 |  |  | addis | Add Immediate Shifted |
| 011111......... \|1111. 011101010. | XO | 1 | 72 | P1 |  | SR | addme[o][.] | Add to Minus One Extended |
| 011111......... 11111.011001010. | XO | I | 73 | P1 |  | SR | addze[0][.] | Add to Zero Extended |
| 011111 $\ldots \ldots . . . . . . . . . . . ~ 0000011100$. | X | 1 | 93 | P1 |  | SR | and[.] | AND |
| 011111.............. 0000111100. | X | I | 94 | P1 |  | SR | andc[.] | AND with Complement |
|  | D | I | 91 | P1 |  | SR | andi. | AND Immediate \& record |
|  | D | I | 91 | P1 |  | SR | andis. | AND Immediate Shifted \& record |
|  | I | I | 38 | P1 |  |  | $\mathrm{b}[1][\mathrm{a}]$ | Branch [\& Link] [Absolute] |
|  | B | 1 | 38 | P1 |  | CT | bc[l] [a] | Branch Conditional [\& Link] [Absolute] |
| 010011.........111.. 1000010000. | XL | I | 39 | P1 |  | CT | bcctr[[] | Branch Conditional to CTR [\& Link] |
| 010011........ 111.. 0000010000. | XL | I | 39 | P1 |  | CT | bclr[[] | Branch Conditional to LR [\& Link] |
| 011111...1. ......... 00000000001 | X | I | 85 | P1 |  |  | cmp | Compare |
|  | D | I | 85 | P1 |  |  | cmpi | Compare Immediate |
| $011111 \ldots . .1 . \ldots . . . . .0000100001$ | X | I | 86 | P1 |  |  | cmpl | Compare Logical |
|  | D | I | 86 | P1 |  |  | cmpli | Compare Logical Immediate |
| 011111......... 111110000011010. | X | I | 95 | P1 |  | SR | cntlzw[.] | Count Leading Zeros Word |
| $010011 \ldots \ldots . \ldots \ldots .01000000011$ | XL | I | 41 | P1 |  |  | crand | CR AND |
| 010011 $\ldots \ldots . \ldots \ldots . . . . . .00100000011$ | XL | I | 42 | P1 |  |  | crandc | CR AND with Complement |
|  | XL | I | 42 | P1 |  |  | creqv | CR Equivalent |
| $010011 \ldots \ldots . \ldots \ldots .00111000011$ | XL | I | 41 | P1 |  |  | crnand | CR NAND |
| $010011 \ldots \ldots . \ldots \ldots .00001000011$ | XL | I | 42 | P1 |  |  | crnor | CR NOR |
|  | XL | I | 41 | P1 |  |  | cror | CR OR |
| 010011............. 01101000011 | XL | I | 42 | P1 |  |  | crorc | CR OR with Complement |
|  | XL | I | 41 | P1 |  |  | crxor | CR XOR |
| 011111 11111 $\ldots . . . . . . . . .11111101101$ | X | 11 | 853 | P1 |  |  | dcbz | Data Cache Block Zero |
|  | X | I | 94 | P1 |  | SR | eqv[.] | Equivalent |
| 011111 $\ldots . .$. $\ldots 111111110011010$ | X | I | 94 | P1 |  | SR | extsh[.] | Extend Sign Hword |
| 111111..... $11111 \ldots . . .0100001000$. | X | , | 151 | P1 |  |  | fabs[.] | Floating Absolute |
| 111111............. 1111110101. | A | I | 153 | P1 |  |  | fadd[.] | Floating Add |
| 111111...11 $\ldots . . . . . .00001000001$ | X | I | 168 | P1 |  |  | fcmpo | Floating Compare Ordered |
| 111111...11 $\ldots . . . . . .00000000001$ | X | I | 168 | P1 |  |  | fcmpu | Floating Compare Unordered |
| 111111.............. 1111110010. | A | I | 154 | P1 |  |  | fdiv[.] | Floating Divide |
| 1111111.... $\ldots . . .1 . . . . . . . . .11101$. | A | I | 158 | P1 |  |  | fmadd[.] | Floating Multiply-Add |
| 111111.... $11111 \ldots \ldots 0001001000$. | X | I | 151 | P1 |  |  | fmr[.] | Floating Move Register |
|  | A | I | 159 | P1 |  |  | fmsub[.] | Floating Multiply-Subtract |
| 111111......... 11111 ..... 11001. | A | I | 154 | P1 |  |  | fmul[.] | Floating Multiply |
| 111111.... $11111 \ldots . .10010001000$. | X | I | 151 | P1 |  |  | fnabs[.] | Floating Negative Absolute Value |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 15 of 17)


Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 16 of 17)

| Instruction $^{\mathbf{1}}$11 11111 $11112 \quad 22222222233$ <br> $012345 \quad 67890$ 12345 $67890 \quad 12345 \quad 678901$ | $\begin{aligned} & \text { \# } \\ & \text { 틍 } \\ & \end{aligned}$ | $\begin{aligned} & \text { ro } \\ & \text { on } \\ & \hline \end{aligned}$ | $\begin{aligned} & 0 \\ & 0 \\ & 0 \end{aligned}$ |  |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | D | I | 91 | P1 |  |  | ori | OR Immediate |
|  | D | I | 92 | P1 |  |  | oris | OR Immediate Shifted |
|  | M | I | 102 | P1 |  | SR | rlwimi[.] | Rotate Left Word Immediate then Mask Insert |
|  | M | I | 101 | P1 |  | SR | rlwinm[.] | Rotate Left Word Immediate then AND with Mask |
|  | M | I | 102 | P1 |  | SR | rlwnm[.] | Rotate Left Word then AND with Mask |
|  | X | I | 106 | P1 |  | SR | slw[.] | Shift Left Word |
| $011111 \ldots . . . . . . . . . .1100011000$. | X | 1 | 107 | P1 |  | SR | sraw[.] | Shift Right Algebraic Word |
| 011111 $\ldots .$. $\ldots . .$. 1100111000. | X | I | 107 | P1 |  | SR | srawi[.] | Shift Right Algebraic Word Immediate |
| 0111111............. 1000011000. | X | 1 | 106 | P1 |  | SR | Srw[.] | Shift Right Word |
|  | D | I | 55 | P1 |  |  | stb | Store Byte |
| 100111......................... | D | I | 55 | P1 |  |  | stbu | Store Byte with Update |
| $011111 \ldots \ldots . . . . . . . . .00111101111$ | X | I | 55 | P1 |  |  | stbux | Store Byte with Update Indexed |
| 011111.............. 00110101111 | X | I | 55 | P1 |  |  | stbx | Store Byte Indexed |
|  | D | I | 147 | P1 |  |  | stfd | Store Floating Double |
| 110111.............. | D | I | 147 | P1 |  |  | stfdu | Store Floating Double with Update |
| 011111............... 10111101111 | X | I | 147 | P1 |  |  | stfdux | Store Floating Double with Update Indexed |
| $011111 \ldots . . . . . . . . . . .10110101111$ | X | I | 147 | P1 |  |  | stfdx | Store Floating Double Indexed |
|  | D | 1 | 146 | P1 |  |  | stfs | Store Floating Single |
| 110101......................... | D | I | 146 | P1 |  |  | stfsu | Store Floating Single with Update |
| 011111..... $\ldots . . . . . . .10101101111$ | X | I | 146 | P1 |  |  | stfsux | Store Floating Single with Update Indexed |
| $011111 \ldots \ldots . . . . . . . .10100101111$ | X | I | 146 | P1 |  |  | stfsx | Store Floating Single Indexed |
| 101100 ..... ..... ..... | D | 1 | 56 | P1 |  |  | sth | Store Hword |
| $011111 \ldots \ldots . . . . . . . . . .11100101101$ | X | I | 61 | P1 |  |  | sthbrx | Store Hword Byte-Reverse Indexed |
| 101101..... ..... ..... ..... ...... | D | I | 56 | P1 |  |  | sthu | Store Hword with Update |
| 011111.............. 01101101111 | X | I | 56 | P1 |  |  | sthux | Store Hword with Update Indexed |
| 011111............. 01100101111 | X | 1 | 56 | P1 |  |  | sthx | Store Hword Indexed |
| 101111....................... | D | I | 63 | P1 |  |  | stmw | Store Multiple Word |
| 011111............. 10110101011 | X | I | 66 | P1 |  |  | stswi | Store String Word Immediate |
| 0111111.... $\ldots \ldots . . . . . .10100101011$ | X | I | 66 | P1 |  |  | stSwX | Store String Word Indexed |
| 100100 ......... ..... ..... ..... | D | I | 57 | P1 |  |  | stw | Store Word |
| 011111.............. 10100101101 | X | I | 61 | P1 |  |  | stwbrx | Store Word Byte-Reverse Indexed |
|  | D | I | 57 | P1 |  |  | stwu | Store Word with Update |
| $011111 \ldots . . . . . . . . . . .00101101111$ | X | I | 57 | P1 |  |  | stwux | Store Word with Update Indexed |
| 011111............. 00100101111 | X | 1 | 57 | P1 |  |  | stwx | Store Word Indexed |
| 011111................ 000001000. | XO | I | 71 | P1 |  | SR | subfc[0][.] | Subtract From Carrying |
| 011111..... .......... 010001000. | XO | 1 | 72 | P1 |  | SR | subfe[0][.] | Subtract From Extended |
| $001000 \ldots \ldots . . . . . . . . . . . . . . . . . . . . ~$ | D | I | 71 | P1 |  | SR | subfic | Subtract From Immediate Carrying |
| 011111......... 11111.011101000. | XO | I | 72 | P1 |  | SR | subfme[0][.] | Subtract From Minus One Extended |
| 011111......... 11111.011001000. | XO | I | 73 | P1 |  | SR | subfze[0][.] | Subtract From Zero Extended |
| $011111111 . .11\|1\| 1\|1\| \mid 10010101101$ | X | 11 | 877 | P1 |  |  | sync | Synchronize |
| 0111111111. 11111.... 01001100101 | X | 111 | 1034 | P1 | H | 64 | tlbie | TLB Invalidate Entry |
| 011111............ 00000001001 | X | I | 89 | P1 |  |  | tw | Trap Word |
|  | D | I | 89 | P1 |  |  | twi | Trap Word Immediate |
|  | X | I | 93 | P1 |  | SR | xor[.] | XOR |
| 011010 ..... ..... ..... ..... ...... | D | , | 92 | P1 |  |  | xori | XOR Immediate |
| 011011..... ..... ..... ..... ...... | D | I | 92 | P1 |  |  | xoris | XOR Immediate Shifted |

Figure 88. Power ISA Instruction Set Sorted by Version (Sheet 17 of 17)

1. Key to Instruction column (primary and extended opcode bits shaded in gray).
| Instruction bit that corresponds to a reserved field, must have a value of 0 , otherwise invalid form

- Instruction bit that corresponds to an operand bit, may have a value of either 0 or 1.
$0 \quad$ Instruction bit having a value 0.
1 Instruction bit having a value 1 .

2. Key to Version column.

| P1 | Instruction introduced in the POWER Architecture. |
| :---: | :--- |
| P2 | Instruction introduced in the POWER2 Architecture. |
| PPC | Instruction introduced in the PowerPC Architecture prior to v2.00. |
| v2.00 | Instruction introduced in the PowerPC Architecture Version 2.00. |
| v2.01 | Instruction introduced in the PowerPC Architecture Version 2.01. |
| v2.02 | Instruction introduced in the PowerPC Architecture Version 2.02. |
| v2.03 | Instruction introduced in the Power ISA Architecture Version 2.03. |
| v2.04 | Instruction introduced in the Power ISA Architecture Version 2.04. |
| v2.05 | Instruction introduced in the Power ISA Architecture Version 2.05. |
| v2.06 | Instruction introduced in the Power ISA Architecture Version 2.06. |
| v2.07 | Instruction introduced in the Power ISA Architecture Version 2.07. |
| v3.0 | Instruction introduced in the Power ISA Architecture Version 3.00. |

3. Key to Privilege column.

P Denotes an instruction that is treated as privileged.
O Denotes an instruction that is treated as privileged or nonprivileged (or hypervisor, for mtspr), depending on the SPR or PMR number.
PI Denotes an instruction that is illegal in privileged state.
H Denotes an instruction that can be executed only in hypervisor state
U Denotes an instruction that can be executed only in ultravisor state
4. Key to Mode Dependency column.

Except as described below and in Section 1.11.3, "Effective Address Calculation", in Book I, all instructions are independent of whether the processor is in 32-bit or 64-bit mode.

[^17]
## Appendix F. Power ISA Instruction Set Sorted by Mnemonic

This appendix lists all the instructions in the Power ISA, sorted by mnemonic.


Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 1 of 17)

| Instruction ${ }^{\mathbf{1}}$ 01111111111222222222233 012345678901234567891234678901 |  | $\begin{aligned} & \text { 믕 } \\ & \text { O} \end{aligned}$ | 告 |  |  | $\left\|\begin{array}{l} r_{0}^{0} \\ 0 \\ 0 \\ \frac{0}{0} \\ 0 \\ \dot{\Sigma} \end{array}\right\|$ |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $011111 \ldots 11 \ldots . . . . .00111000001$ | X | 1 | 88 | v3.0 |  |  | cmpeqb | Compare Equal Byte |
| 001011 ...1. ..... ..... ..... ...... | D | I | 85 | P1 |  |  | cmpi | Compare Immediate |
| $011111 \ldots 1 . \ldots \ldots \ldots 0.00001000001$ | X | 1 | 86 | P1 |  |  | cmpl | Compare Logical |
| $001010 \ldots 1 . \ldots . . . . . . . .$. | D | 1 | 86 | P1 |  |  | cmpli | Compare Logical Immediate |
| $011111 . . .1 . \ldots . . . . . .000110000001$ | X | 1 | 87 | v3.0 |  |  | cmprb | Compare Ranged Byte |
| $011111 \ldots . . . . . .111110000111010$. | X | 1 | 98 | PPC |  | SR | cntzd[[] | Count Leading Zeros Dword |
| $011111 \ldots . . . . .111110000011010$. | X | 1 | 95 | P1 |  | SR | cntzw[.] | Count Leading Zeros Word |
|  | X | 1 | 98 | v3.0 |  |  | cntzd[].] | Count Trailing Zeros Dword |
| $011111 \ldots . . . . . .111111000011010$. | X | 1 | 95 | v3.0 |  |  | cnttzw[.] | Count Trailing Zeros Word |
| 011111 \|111. ......... 11000001101 | X | 11 | 858 | v3.0 |  |  | copy | Copy |
| 011111 \|1|1| ||11| |1|1| 11010001101 | X | 11 | 860 | v3.0 |  |  | Cp_abort | CP_Abort |
|  | XL | 1 | 41 | P1 |  |  | crand | CR AND |
| $010011 \ldots . . . . . . . . . . .00100000011$ | XL | I | 42 | P1 |  |  | crandc | CR AND with Complement |
| $010011 \ldots \ldots \ldots \ldots . . .101001000011$ | XL | 1 | 42 | P1 |  |  | creqv | CR Equivalent |
| $010011 \ldots . . . . . . . . . . .000111000011$ | XL | 1 | 41 | P1 |  |  | crnand | CR NAND |
| $010011 \ldots . . . . . . . . . .00001000011$ | XL | I | 42 | P1 |  |  | crnor | CR NOR |
|  | XL | 1 | 41 | P1 |  |  | cror | CR OR |
| $010011 \ldots . . . . . . . . . . .01101000011$ | XL | 1 | 42 | P1 |  |  | crorc | CR OR with Complement |
| 010011 .............. 00110000011 | XL | I | 41 | P1 |  |  | crxor | CR XOR |
| 111011.............. 0000000010. | X | I | 195 | v2.05 |  |  | dadd[.] | DFP Add |
| 111111............... 0000000010. | X | 1 | 195 | v2.05 |  |  | daddq[] | DFP Add Quad |
| 011111 ..... \|11.. |111| $1011110011 /$ | X | 1 | 79 | v3.0 |  |  | darn | Deliver A Random Number |
| 011111 \|11....... .... 00010101101 | X | 11 | 854 | PPC |  |  | dcbf | Data Cache Block Flush |
| $01111111111 \ldots \ldots . . . .00001101101$ | X | 11 | 853 | PPC |  |  | dcbst | Data Cache Block Store |
|  | X | 11 | 851 | PPC |  |  | dcbt | Data Cache Block Touch |
| 011111 .............. 00111101101 | X | 11 | 852 | PPC |  |  | dcbist | Data Cache Block Touch for Store |
| 011111 \|1111 .......... 11111101101 | X | 11 | 853 | P1 |  |  | dcbz | Data Cache Block Zero |
| $111011 \ldots . . \mid 1111 \ldots . .1100100010$. | X | 1 | 217 | v2.06 |  |  | dcffix[.] | DFP Convert From Fixed |
| 1111111.... $11111 \ldots . . .1100100010$. | X | 1 | 217 | v2.05 |  |  | dcffixq[.] | DFP Convert From Fixed Quad |
| 111011 ...11 ......... 00100000101 | X | 1 | 201 | v2.05 |  |  | dcmpo | DFP Compare Ordered |
| 111111 ...11 ......... 00100000101 | X | I | 201 | v2.05 |  |  | dcmpoq | DFP Compare Ordered Quad |
| 111011 ...11 ......... 10100000101 | X | 1 | 200 | v2.05 |  |  | dcmpu | DFP Compare Unordered |
| 111111 ...11 ......... 10100000101 | X | 1 | 200 | v2.05 |  |  | dcmpuq | DFP Compare Unordered Quad |
| 111011.... $11111 \ldots \ldots 0100000010$. | X | I | 215 | v2.05 |  |  | dctdp[.] | DFP Convert To DFP Long |
| $111011 \ldots . .11111 \ldots . .10100100010$. | X | I | 217 | v2.05 |  |  | dctixi[.] | DFP Convert To Fixed |
| 111111.... $11111 \ldots . .0100100010$. | X | 1 | 217 | v2.05 |  |  | dctixal.] | DFP Convert To Fixed Quad |
|  | X | I | 215 | v2.05 |  |  | dctapq[].] | DFP Convert To DFP Extended |
| 1111011 .......111 .... 0101000010. | X | 1 | 219 | v2.05 |  |  | ddedpd[.] | DFP Decode DPD To BCD |
| 111111.......111 .... 0101000010. | X | 1 | 219 | v2.05 |  |  | ddedpdq[.] | DFP Decode DPD To BCD Quad |
| 111011.............. 1000100010. | X | 1 | 198 | v2.05 |  |  | ddiv[.] | DFP Divide |
| 111111.......... ..... 1000100010. | X | 1 | 198 | v2.05 |  |  | ddivg[.] | DFP Divide Quad |
| 111011 ..... .1111 ..... 1101000010. | X | 1 | 219 | v2.05 |  |  | denbcd[.] | DFP Encode BCD To DPD |
| 1111111..... $1111 \ldots \ldots .1101000010$. | X | 1 | 219 | v2.05 |  |  | denbcdq[.] | DFP Encode BCD To DPD Quad |
| 111011 .......... .... 1101100010. | X | 1 | 220 | v2.05 |  |  | diex[] | DFP Insert Exponent |
| 111111..... ..... ..... 1101100010. | X | 1 | 220 | v2.05 |  |  | diexa[.] | DFP Insert Exponent Quad |
|  | XO | 1 | 82 | PPC |  | SR | divd[[0][] | Divide Dword |
|  | XO | 1 | 83 | v2.06 |  | SR | divde[0][.] | Divide Dword Extended |
|  | XO | I | 83 | v2.06 |  | SR | divdeu[0][] | Divide Dword Extended Unsigned |
| 011111 ..... ..... ..... . 1111001001. | XO | I | 82 | PPC |  | SR | divdu[0][]] | Divide Dword Unsigned |
|  | XO | 1 | 75 | PPC |  | SR | diww[0][] | Divide Word |
| 011111..... ..... ..... . 110101011. | XO | I | 77 | v2.06 |  | SR | diwwe[0][] | Divide Word Extended |
|  | XO | I | 77 | v2.06 |  | SR | diwweu[0][.] | Divide Word Extended Unsigned |
|  | XO | 1 | 75 | PPC |  | SR | diwwu[0][]] | Divide Word Unsigned |
| 111011.............. 0000100010. | X | I | 197 | v2.05 |  |  | dmul[.] | DFP Multiply |
| $\begin{array}{lllllll}111111 & \ldots . . . & \ldots . . & 0000100010 .\end{array}$ | X | 1 | 197 | v2.05 |  |  | dmulq[.] | DFP Multiply Quad |
| $111011 \ldots . . . . . . . . . . . . . . .00000011$. | Z23 | I | 206 | v2.05 |  |  | dqua[.] | DFP Quantize |

Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 2 of 17)
I

| Instruction $^{\mathbf{1}}$ 01234567890111111111222222222233 123489612345678901 |  | $\begin{aligned} & \text { ro } \\ & \text { on } \end{aligned}$ | $\begin{aligned} & 0 \\ & \underset{\sim}{0} \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \stackrel{0}{0} \\ & \text { N } \\ & \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 111011................. 01000011. | Z23 | I | 205 | v2.05 |  |  | dquai[.] | DFP Quantize Immediate |
| 111111............. . 01000011. | Z23 | I | 205 | v2.05 |  |  | dquaiq[.] | DFP Quantize Immediate Quad |
| 111111................ 00000011. | Z23 | I | 206 | v2.05 |  |  | dquaq[.] | DFP Quantize Quad |
| 111111.... 11111.... 1100000010. | X | I | 216 | v2.05 |  |  | drdpq[.] | DFP Round To DFP Long |
| 111011.... 1111. ..... . 111100011. | Z23 | I | 213 | v2.05 |  |  | drintn[.] | DFP Round To FP Integer Without Inexact |
| 111111.... 1111. .... . 11100011. | Z23 | 1 | 213 | v2.05 |  |  | drintnq[.] | DFP Round To FP Integer Without Inexact Quad |
| 111011.... 1111. ...... 01100011. | Z23 | 1 | 211 | v2.05 |  |  | drintX[.] | DFP Round To FP Integer With Inexact |
| 111111..... 1111. ..... . 011100011. | Z23 | I | 211 | v2.05 |  |  | drintxq[] | DFP Round To FP Integer With Inexact Quad |
| 111011................. 00100011. | Z23 | 1 | 208 | v2.05 |  |  | drrnd[.] | DFP Reround |
| 111111................ 00100011. | Z23 | I | 208 | v2.05 |  |  | drrndq[] | DFP Reround Quad |
| 111011.... 11111 .... 1100000010. | X | I | 216 | v2.05 |  |  | drsp[.] | DFP Round To DFP Short |
| 111011............. . 001000010. | Z22 | I | 222 | v2.05 |  |  | dscli[.] | DFP Shift Significand Left Immediate |
| 111111............. . 001000010. | Z22 | 1 | 222 | v2.05 |  |  | dscliq[.] | DFP Shift Significand Left Immediate Quad |
| 111011............. . 001100010. | Z22 | I | 222 | v2.05 |  |  | dscri[.] | DFP Shift Significand Right Immediate |
| 111111.............. 001100010. | Z22 | I | 222 | v2.05 |  |  | dscriq[.] | DFP Shift Significand Right Immediate Quad |
| 111011............. 1000000010. | X | I | 195 | v2.05 |  |  | dsub[.] | DFP Subtract |
| 111111 $\ldots \ldots \ldots \ldots \ldots .1000000010$. | X | 1 | 195 | v2.05 |  |  | dsubq[.] | DFP Subtract Quad |
| 111011 $\ldots .11 \ldots . . . . . . . .0110000101$ | Z22 | 1 | 202 | v2.05 |  |  | dtstdc | DFP Test Data Class |
| 111111...11 .......... 0110000101 | Z22 | 1 | 202 | v2.05 |  |  | dtstdcq | DFP Test Data Class Quad |
| 111011...11 .......... . 0111000101 | Z22 | I | 202 | v2.05 |  |  | dtstdg | DFP Test Data Group |
| 111111...11 ..... ..... . 0111000101 | Z22 | 1 | 202 | v2.05 |  |  | dtstdgq | DFP Test Data Group Quad |
| 111011...11 .... ..... 00101000101 | X | 1 | 203 | v2.05 |  |  | dtstex | DFP Test Exponent |
| 111111...11 $\ldots . . . \ldots . .00101000101$ | X | 1 | 203 | v2.05 |  |  | dtstexq | DFP Test Exponent Quad |
| 111011...11 ......... 10101 000101 | X | I | 204 | v2.05 |  |  | dtstsf | DFP Test Significance |
| 111011 $\ldots 11 \ldots \ldots \ldots .10101000111$ | X | 1 | 204 | v3.0 |  |  | dtstsfi | DFP Test Significance Immediate |
| 111111 $\ldots .11 \ldots . . . . . .10101000111$ | X | 1 | 204 | v3.0 |  |  | dtstsfiq | DFP Test Significance Immediate Quad |
| 111111 $\ldots .11 \ldots . . . . . . .10101000101$ | X | I | 204 | v2.05 |  |  | dtstsfq | DFP Test Significance Quad |
| 111011 $\ldots . .11111 \ldots \ldots 0101100010$. | X | I | 220 | v2.05 |  |  | dxex[.] | DFP Extract Exponent |
|  | X | I | 220 | v2.05 |  |  | dxexq[.] | DFP Extract Exponent Quad |
| 01111111111111111111111010101101 | X | 11 | 879 | PPC |  |  | eieio | Enforce In-order Execution of I/O |
|  | X | 1 | 94 | P1 |  | SR | eqv[.] | Equivalent |
|  | X | 1 | 94 | PPC |  | SR | extsb[.] | Extend Sign Byte |
|  | X | I | 94 | P1 |  | SR | extsh[.] | Extend Sign Hword |
|  | X | I | 98 | PPC |  | SR | extsw[.] | Extend Sign Word |
| 011111..... $\ldots$... $\ldots$.... $110111101 .$. | XS | I | 109 | v3.0 |  |  | extswsli[.] | Extend Sign Word \& Shift Left Immediate |
| 111111 $\ldots \ldots .11111 \ldots \ldots 0100001000$. | X | 1 | 151 | P1 |  |  | fabs[.] | Floating Absolute |
| 111111.... .... .... 1111110101. | A | 1 | 153 | P1 |  |  | fadd[.] | Floating Add |
| 111011.............. 1111110101. | A | 1 | 153 | PPC |  |  | fadds[.] | Floating Add Single |
| 111111.... $11111 \ldots \ldots 1101001110$. | X | I | 164 | PPC |  |  | fcfid[.] | Floating Convert From Integer Dword |
| 111011 .... 11111 .... 1101001110. | X | I | 165 | v2.06 |  |  | fcfids[.] | Floating Convert From Integer Dword Single |
| 111111.... $11111 \ldots \ldots 1111001110$. | X | 1 | 165 | v2.06 |  |  | fcfidu[.] | Floating Convert From Integer Dword Unsigned |
| 111011 $\ldots . . .11111 \ldots . .1111001110$. | X | I | 166 | v2.06 |  |  | fcfidus[.] | Floating Convert From Integer Dword Unsigned Single |
| 111111...11 .... .... 0000100001 | X | I | 168 | P1 |  |  | fcmpo | Floating Compare Ordered |
| 111111...11 ........ 0000000001 | X | I | 168 | P1 |  |  | fcmpu | Floating Compare Unordered |
| 111111............. 0000001000. | X | 1 | 151 | v2.05 |  |  | fcpsgn[.] | Floating Copy Sign |
| 111111.... $11111 \ldots \ldots 1100101110$. | X | I | 160 | PPC |  |  | fctid[.] | Floating Convert To Integer Dword |
| 111111 $\ldots . . .11111 \ldots . .1110101110$. | X | 1 | 161 | v2.06 |  |  | fctidu[.] | Floating Convert To Integer Dword Unsigned |
|  | X | 1 | 162 | v2.06 |  |  | fctiduz[.] | Floating Convert To Integer Dword Unsigned truncate |
| 111111.... $11111 \ldots . .1100101111$. | X | I | 161 | PPC |  |  | fctidz[.] | Floating Convert To Integer Dword truncate |
| 111111.... $11111 \ldots . . .0000001110$. | X | 1 | 162 | P2 |  |  | fctiw[.] | Floating Convert To Integer Word |
| 111111 $\ldots . . .11111 \ldots . . .0010001110$. | X | 1 | 163 | v2.06 |  |  | fctiwu[.] | Floating Convert To Integer Word Unsigned |
| 111111..... $11111 \ldots \ldots 0010001111$. | X | 1 | 164 | v2.06 |  |  | fctiwuz[.] | Floating Convert To Integer Word Unsigned truncate |
| 111111.... $11111 \ldots . .0000001111$. | X | 1 | 163 | P2 |  |  | fctiwz[.] | Floating Convert To Integer Word truncate |
| 111111............. 1111110010. | A | I | 154 | P1 |  |  | fdiv[.] | Floating Divide |
| 111011 $\ldots \ldots . . . . . . . . . .1111110010$. | A | I | 154 | PPC |  |  | fdivs[.] | Floating Divide Single |
| 111111..... ..... ..... ..... 11101. | A | I | 158 | P1 |  |  | fmadd[.] | Floating Multiply-Add |

Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 3 of 17)


Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 4 of 17)

|  |  | Y O 0 0 | $\begin{aligned} & 0 \\ & 0 \\ & 00 \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \stackrel{0}{n} \\ & \text { N } \\ & \hline \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 011111............... 11010101111 | X | I | 144 | v2.05 |  |  | Ifiwax | Load Floating as Integer Word Algebraic Indexed |
| 011111..... .... ..... 11011 101111 | X | I | 144 | v2.06 |  |  | Ifiwzx | Load Floating as Integer Word \& Zero Indexed |
| $110000 \ldots . . . . . . . . . . . . . . . . . . . . . . ~$ | D | 1 | 142 | P1 |  |  | Ifs | Load Floating Single |
| 110001......................... | D | I | 142 | P1 |  |  | Ifsu | Load Floating Single with Update |
| 011111..... .... ..... 10001101111 | X | I | 142 | P1 |  |  | Ifsux | Load Floating Single with Update Indexed |
| 011111............. 10000101111 | X | I | 142 | P1 |  |  | Ifsx | Load Floating Single Indexed |
| 101010 ........................ | D | I | 50 | P1 |  |  | Iha | Load Hword Algebraic |
| 011111.............. 0001110100. | X | 11 | 869 | v2.06 |  |  | Iharx | Load Hword And Reserve Indexed Xform |
| 101011............... | D | I | 50 | P1 |  |  | Ihau | Load Hword Algebraic with Update |
| 011111..... .......... 01011101111 | X | I | 50 | P1 |  |  | Ihaux | Load Hword Algebraic with Update Indexed |
| 011111.............. 01010101111 | X | I | 50 | P1 |  |  | Ihax | Load Hword Algebraic Indexed |
| 011111 $\ldots \ldots . . . . . . . . . . .11000101101$ | X | I | 61 | P1 |  |  | Ihbrx | Load Hword Byte-Reverse Indexed |
| 101000 ..... ..... ..... ..... ...... | D | I | 49 | P1 |  |  | Ihz | Load Hword \& Zero |
| 011111 $\ldots \ldots . . . . . . . . . . .11001101011$ | X | 111 | 966 | v2.05 | H |  | Ihzcix | Load Hword \& Zero Caching Inhibited Indexed |
| 101001......................... | D | I | 49 | P1 |  |  | Ihzu | Load Hword \& Zero with Update |
| 011111............... 01001101111 | X | I | 49 | P1 |  |  | Ihzux | Load Hword \& Zero with Update Indexed |
| 011111............. 0100010111 | X | I | 49 | P1 |  |  | Ihzx | Load Hword \& Zero Indexed |
|  | D | I | 63 | P1 |  |  | Imw | Load Multiple Word |
| 111000 ..... ..... .......... . | DQ | I | 59 | v2.03 |  |  | Iq | Load Qword |
|  | X | 11 | 875 | v2.07 |  |  | Iqarx | Load Qword And Reserve Indexed |
|  | X | I | 65 | P1 |  |  | Iswi | Load String Word Immediate |
| 011111......... ..... 10000101011 | X | 1 | 65 | P1 |  |  | Iswx | Load String Word Indexed |
| $011111 \ldots \ldots . . . . . . . . . . .00000001111$ | X | I | 244 | v2.03 |  |  | Ivebx | Load Vector Element Byte Indexed |
| $011111 \ldots . . . . . . . . . . . .00001001111$ | X | I | 244 | v2.03 |  |  | Ivehx | Load Vector Element Hword Indexed |
|  | X | I | 245 | v2.03 |  |  | Ivewx | Load Vector Element Word Indexed |
| $011111 \ldots . . . . . . . . . .00000001101$ | X | I | 249 | v2.03 |  |  | \|vsl | Load Vector for Shift Left |
| $011111 \ldots \ldots . \ldots \ldots .00001001101$ | X | I | 249 | v2.03 |  |  | Ivsr | Load Vector for Shift Right |
| $011111 \ldots \ldots . . . . . . . . . . .00011001111$ | X | I | 245 | v2.03 |  |  | IVX | Load Vector Indexed |
|  | X | I | 245 | v2.03 |  |  | \|vx| | Load Vector Indexed Last |
|  | DS | I | 52 | PPC |  |  | Iwa | Load Word Algebraic |
| 011111.............. 00000101001 | X | 11 | 869 | PPC |  |  | Iwarx | Load Word \& Reserve Indexed |
| $011111 \ldots \ldots . \ldots \ldots .10010001101$ | X | 11 | 864 | v3.0 |  |  | Iwat | Load Word ATomic |
| 011111............... $0101110101 /$ | X | I | 52 | PPC |  |  | Iwaux | Load Word Algebraic with Update Indexed |
| 011111......... .... $0101010101 /$ | X | I | 52 | PPC |  |  | Iwax | Load Word Algebraic Indexed |
| 011111.............. 10000101101 | X | I | 61 | P1 |  |  | Iwbrx | Load Word Byte-Reverse Indexed |
|  | D | I | 51 | P1 |  |  | IWz | Load Word \& Zero |
|  | X | 111 | 966 | v2.05 | H |  | Iwzcix | Load Word \& Zero Caching Inhibited Indexed |
| 100001......................... | D | I | 51 | P1 |  |  | Iwzu | Load Word \& Zero with Update |
| 011111............. 00001101111 | X | I | 51 | P1 |  |  | IWzux | Load Word \& Zero with Update Indexed |
| $011111 \ldots \ldots . . . . . . . . .00000101111$ | X | I | 51 | P1 |  |  | IWZX | Load Word \& Zero Indexed |
|  | DS | I | 481 | v3.0 |  |  | Ixsd | Load VSX Scalar Dword |
| 011111 $\ldots . . . . . . . . . . . . .1001001100$. | XX1 | I | 481 | v2.06 |  |  | Ixsdx | Load VSX Scalar Dword Indexed |
| $011111 \ldots \ldots . . . . . . . . .1100001101$. | XX1 | I | 483 | v3.0 |  |  | Ixsibzx | Load VSX Scalar as Integer Byte \& Zero Indexed |
|  | XX1 | 1 | 483 | v3.0 |  |  | Ixsihzx | Load VSX Scalar as Integer Hword \& Zero Indexed |
| 011111............. 0001001100. | XX1 | I | 484 | v2.07 |  |  | Ixsiwax | Load VSX Scalar as Integer Word Algebraic Indexed |
|  | XX1 | I | 485 | v2.07 |  |  | Ixsiwzx | Load VSX Scalar as Integer Word \& Zero Indexed |
|  | DS | I | 486 | v3.0 |  |  | Ixssp | Load VSX Scalar Single |
| $011111 \ldots \ldots . . . . . . . . . .1000001100$. | XX1 | I | 486 | v2.07 |  |  | Ixsspx | Load VSX Scalar SP Indexed |
| 111101 ..... ..... ..... ..... ... 001 | DQ | I | 493 | v3.0 |  |  | Ixv | Load VSX Vector |
| 01111 $\ldots . . . . . . . . . .11011101100 . ~$ | XX1 | I | 488 | v3.0 |  |  | Ixvb16x | Load VSX Vector Byte*16 Indexed |
| 011111 $\ldots \ldots . . . . . . . . . . . .1101001100$. | XX1 | I | 489 | v2.06 |  |  | Ixvd2x | Load VSX Vector Dword*2 Indexed |
| 011111..... $\ldots \ldots . . . . . .0101001100$. | XX1 | I | 495 | v2.06 |  |  | Ixvdsx | Load VSX Vector Dword \& Splat Indexed |
| 011111 $\ldots . .$. $\ldots . .$. 1100101100. | XX1 | 1 | 496 | v3.0 |  |  | Ixvh8x | Load VSX Vector Hword*8 Indexed |
|  | XX1 | 1 | 490 | v3.0 |  |  | \|xv| | Load VSX Vector with Length |
| 011111 $\ldots \ldots . . . . . . . . . .0100101101$. | XX1 | I | 492 | v3.0 |  |  | Ixv\|l | Load VSX Vector Left-justified with Length |
| $011111 \ldots . . . . . . . . . . .11100001100$. | XX1 | I | 497 | v2.06 |  |  | Ixww4x | Load VSX Vector Word*4 Indexed |

Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 5 of 17)

| Instruction ${ }^{\mathbf{1}}$ 1111111111222222222233 01234567890123456789012345678901 | $\begin{aligned} & \text { だ } \\ & \text { ت゙ㄹ } \end{aligned}$ | $\begin{aligned} & \text { ᄃ } \\ & \text { O } \\ & \hline \end{aligned}$ | $\begin{aligned} & \mathbb{\otimes} \\ & \stackrel{\pi}{\circ} \end{aligned}$ |  |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | XX1 | 1 | 498 | v3．0 |  |  | Ixwsx | Load VSX Vector Word \＆Splat Indexed |
| 011111．．．．．．．．．．．．．． 0100001100. | XX1 | I | 493 | v3．0 |  |  | lxvx | Load VSX Vector Indexed |
| $000100 \ldots \ldots . . . . . . . . . . . .110000$ | VA | 1 | 81 | v3．0 |  |  | maddhd | Multiply－Add High Dword |
|  | VA | 1 | 81 | v3．0 |  |  | maddhdu | Multiply－Add High Dword Unsigned |
| $000100 \ldots \ldots . . . . . . . . . . . . .110011$ | VA | 1 | 81 | v3．0 |  |  | maddld | Multiply－Add Low Dword |
| $010011 \ldots 11 \ldots 111111100000000001$ | XL | 1 | 42 | P1 |  |  | morf | Move CR Field |
| 111111 ．．．11 ．．．11 11111 00010000001 | X | 1 | 171 | P1 |  |  | mcrfs | Move To CR from FPSCR |
| $011111 \ldots 11111111111110010000001$ | X | 1 | 119 | v3．0 |  |  | mcrix | Move XER to CR Extended |
| 011111 ．．．．．．．．．．．．．．． 01001011101 | X | I | 44 | v2．07 |  |  | mfbhrbe | Move From BHRB |
| $011111 \ldots . .01111 \mid 111100000100111$ | XFX | 1 | 121 | P1 |  |  | mfcr | Move From CR |
| 111111 ．．．．｜｜11｜｜1｜1｜ 1001000111. | X | 1 | 171 | P1 |  |  | mffs［］ | Move From FPSCR |
| $011111 \ldots . .1\|1\| 1\|1\| 1\|0001010011\|$ | X | 111 | 979 | P1 | P |  | mfmsr | Move From MSR |
| $011111 \ldots . .1 \ldots . .10 .100000100111$ | XFX | I | 121 | v2．01 |  |  | mfocrf | Move From One CR Field |
| 011111 ．．．．．．．．．．．．．． 01010100111 | X | $\begin{aligned} & 11 \\ & 111 \\ & \hline \end{aligned}$ | $\begin{aligned} & 118 \\ & 975 \end{aligned}$ | P1 | 0 |  | mfspr | Move From SPR |
| 011111 ．．．．．．．．．．．．．．． $0101110011 /$ | X | 11 | 902 | PPC |  |  | mftb | Move From Time Base |
| $000100 \ldots . .111111111111000000100$ | VX | 1 | 364 | v2．03 |  |  | mfvscr | Move From VSCR |
| $011111 \ldots . . . . . . .1 \mid 1110000110011$. | XX1 | 1 | 111 | v2．07 |  |  | mfvsrd | Move From VSR Dword |
| $011111 \ldots . . . . . .111110100110011$. | XX1 | I | 111 | v3．0 |  |  | mfvsrld | Move From VSR Lower Dword |
| $011111 \ldots . . . . . .111110001110011$. | XX1 | 1 | 112 | v2．07 |  |  | mfvsrwz | Move From VSR Word \＆Zero |
| $011111 \ldots . . . . . . . . . . .11000010011$ | X | 1 | 84 | v3．0 |  |  | modsd | Modulo Signed Dword |
|  | X | I | 76 | v3．0 |  |  | modsw | Modulo Signed Word |
| $011111 \ldots . . . . . . . . . . .01000010011$ | X | 1 | 84 | v3．0 |  |  | modud | Modulo Unsigned Dword |
| 011111 ．．．．．．．．．．．．．． 01000010111 | X | 1 | 76 | v3．0 |  |  | moduw | Modulo Unsigned Word |
| $0111111111111111 . . .00111011101$ | X | 111 | 1124 | v2．07 | H |  | msgclr | Message Clear |
| 011111 ｜1111｜1111．．．．． 00101011101 | X | 111 | 1126 | v2．07 | P |  | msgclip | Message Clear Privileged |
| 011111 ｜1111｜1111 ．．．． 00110011101 | X | 111 | 1123 | v2．07 | H |  | msgsnd | Message Send |
| 0111111111111111 ．．．． 00100011101 | X | 111 | 1125 | v2．07 | P |  | msgsndp | Message Send Privileged |
| 011111 ｜｜｜｜｜｜｜1｜｜｜1｜｜｜ 11011101101 | X | 111 | 1126 | v3．0 | H |  | msgsync | Message Synchronize |
| $011111 \ldots . .00 . . . . . .100100100001$ | XFX | 1 | 120 | P1 |  |  | mtcrf | Move To CR Fields |
| 111111．．．．11111｜1111 0001000110. | X | I | 173 | P1 |  |  | mtsb0［．］ | Move To FPSCR Bit 0 |
| $1111111 . . .11111111110000100110$. | X | 1 | 173 | P1 |  |  | mtsb1［．］ | Move To FPSCR Bit 1 |
| 111111．．．．．．．．．．．．．．． 1011000111. | XFL | I | 172 | P1 |  |  | mitsf［］．］ | Move To FPSCR Fields |
| 111111 ．．．11｜111．．．．1 0010000110. | X | 1 | 172 | P1 |  |  | mitsfi［］ | Move To FPSCR Field Immediate |
| $011111 \ldots . .1111 .1111100100100101$ | X | 111 | 977 | P1 | P |  | mtmsr | Move To MSR |
| $011111 \ldots . .1\|11.1\| 11100101100101$ | X | 111 | 978 | PPC | P |  | mtmsrd | Move To MSR Dword |
| 011111 $\ldots . . .1 \ldots . . . . .10 .100100100001$ | XFX | － | 120 | v2．01 |  |  | mtocif | Move To One CR Field |
| 011111 ．．．．．．．．．．．．． $0111010011 /$ | X | $\begin{aligned} & 1 \quad 1 \\ & \prime \prime \prime \end{aligned}$ | $\begin{aligned} & 116 \\ & \hline 974 \end{aligned}$ | P1 | 0 |  | mtspr | Move To SPR |
| $000100 \mid 111111111 \ldots . . .11001000100$ | VX | 1 | 364 | v2．03 |  |  | mtvscr | Move To VSCR |
| $011111 \ldots . . . . . . .111110010110011$. | XX1 | 1 | 113 | v2．07 |  |  | mtvsrd | Move To VSR Dword |
|  | XX1 | I | 114 | v3．0 |  |  | mtvsrdd | Move To VSR Double Dword |
| $011111 \ldots . . . . . . .111110011010011$. | XX1 | － | 113 | v2．07 |  |  | mtvsrwa | Move To VSR Word Algebraic |
| $011111 \ldots . . . . . . .111110110010011$. | XX1 | 1 | 115 | v3．0 |  |  | mtvs＇ws | Move To VSR Word \＆Splat |
| $011111 \ldots . . . . . .111110011110011$. | XX1 |  | 114 | v2．07 |  |  | mtvs＇wz | Move To VSR Word \＆Zero |
| 011111 ．．．．．．．．．．．．．．． 1001001001. | XO | I | 80 | PPC |  | SR | mulhd［．］ | Multiply High Dword |
|  | XO | － | 80 | PPC |  | SR | mulhdu［］ | Multiply High Dword Unsigned |
|  | XO | 1 | 74 | PPC |  | SR | mulhw［］ | Multiply High Word |
|  | XO | 1 | 74 | PPC |  | SR | mulhwu［］ | Multiply High Word Unsigned |
|  | XO | 1 | 80 | PPC |  | SR | mulld［0］［．］ | Multiply Low Dword |
|  | D | 1 | 74 | P1 |  |  | mulli | Multiply Low Immediate |
|  | XO | ＋ | 74 | P1 |  | SR | mullw［0］［］ | Multiply Low Word |
| $011111 . . . . . . . . . . . . . .0111011100$. | X | I | 93 | P1 |  | SR | nand［］ | NAND |
| $011111 \ldots \ldots . . .11111 .001101000$. | XO | 1 | 73 | P1 |  | SR | neg［0］［．］ | Negate |
|  | X | ， | 94 | P1 |  | SR | nor［．］ | NOR |
| $011111 \ldots \ldots . . . . . . . . .0110111100$. | X | 1 | 93 | P1 |  | SR | or［．］ | OR |
|  | X | 1 | 94 | P1 |  | SR | orc［．］ | OR with Complement |

Figure 89．Power ISA Instruction Set Sorted by Mnemonic（Sheet 6 of 17）

| Instruction $^{\mathbf{1}}$ <br> $012345 \quad 67890111111111222222 \quad 222233$ <br> $12345 \quad 6789012345678901$ |  | $\begin{aligned} & \text { 두 } \\ & \text { ob } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & 0 \\ & 0 \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \frac{0}{n} \\ & \text { No } \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $011000 \ldots . . . . . . . . . . . . . . . . . . . . . . .$. | D | । | 91 | P1 |  |  | ori | OR Immediate |
|  | D | I | 92 | P1 |  |  | oris | OR Immediate Shifted |
| $011111111 . \ldots . . . . . .1110000110$. | X | 11 | 859 | v3.0 |  |  | paste[.] | Paste |
| $011111 \ldots . . . . . .1 \mid 11100011110101$ | X | I | 96 | v2.02 |  |  | popcntb | Population Count Byte |
| 011111..... .... 111\|| 01111110101 | X | 1 | 98 | v2.06 |  |  | popentd | Population Count Dword |
| $011111 \ldots . . . . . .1111101011110101$ | X | I | 96 | v2.06 |  |  | popentw | Population Count Words |
| $011111 \ldots \ldots . . .1111100101110101$ | X | I | 97 | v2.05 |  |  | prtyd | Parity Dword |
| 011111..... $\ldots . .11111100100110101$ | X | 1 | 97 | v2.05 |  |  | prtyw | Parity Word |
| 01001111111111111111.00100100101 | XL | 11 | 909 | v2.07 |  |  | rfebb | Return from Event Based Branch |
| 010011 \|1111 11111 11111 00000100101 | XL | 111 | 954 | PPC | P |  | rfid | Return from Interrupt Dword |
| 010011 \|1111 |1111 11111 00010100101 | XL | 111 | 953 | v3.0 | P |  | rfscv | Return From System Call Vectored |
|  | MDS | 1 | 103 | PPC |  | SR | rldcl[.] | Rotate Left Dword then Clear Left |
|  | MDS | I | 103 | PPC |  | SR | rldcr[.] | Rotate Left Dword then Clear Right |
|  | MD | I | 104 | PPC |  | SR | rldic[.] | Rotate Left Dword Immediate then Clear |
|  | MD | I | 104 | PPC |  | SR | rldicl[.] | Rotate Left Dword Immediate then Clear Left |
| 011110 ..... ..... ..... ..... . 001.. | MD | I | 105 | PPC |  | SR | rldicr[.] | Rotate Left Dword Immediate then Clear Right |
|  | MD | I | 105 | PPC |  | SR | rldimi[.] | Rotate Left Dword Immediate then Mask Insert |
|  | M | 1 | 102 | P1 |  | SR | rlwimi[.] | Rotate Left Word Immediate then Mask Insert |
| 010101 ..... ..... ..... ..... ...... | M | I | 101 | P1 |  | SR | rlwinm[.] | Rotate Left Word Immediate then AND with Mask |
|  | M | 1 | 102 | P1 |  | SR | rlwnm[.] | Rotate Left Word then AND with Mask |
|  | SC | 1 | 43 | PPC |  |  | SC | System Call |
| $01000111111111111111 . \ldots . .11101$ | SC | I | 43 | v3.0 |  |  | SCV | System Call Vectored |
| $011111 \ldots \ldots . .111111100100000001$ | X | 1 | 121 | v3.0 |  |  | setb | Set Boolean |
| 011111..... $11111 \ldots \ldots 11110100111$ | X | 111 | 1031 | v2.05 | P | SR | slbfee. | SLB Find Entry ESID \& record |
| $01111111 \ldots 1\|1111111\| 101111100101$ | X | 111 | 1027 | PPC | P |  | slbia | SLB Invalidate All |
|  | X | III | 1024 | PPC | P |  | slbie | SLB Invalidate Entry |
| 011111 $\ldots \ldots .11111 \ldots \ldots 01110100101$ | X | III | 1025 | v3.0 | P |  | slbieg | SLB Invalidate Entry Global |
| 011111 $\ldots$.... \||111 $\ldots$.... 11100100111 | X | 111 | 1030 | v2.00 | P |  | slbmfee | SLB Move From Entry ESID |
|  | X | 111 | 1030 | v2.00 | P |  | slbmfev | SLB Move From Entry VSID |
| 011111 $\ldots \ldots .11111 \ldots \ldots .01100100101$ | X | 111 | 1029 | v2.00 | P |  | slbmte | SLB Move To Entry |
| 011111 11111 \|1111 11111 01010100101 | X | 111 | 1031 | v3.0 | P |  | slbsync | SLB Synchronize |
|  | X | 1 | 108 | PPC |  | SR | sld[.] | Shift Left Dword |
| 011111............. 0000011000. | X | 1 | 106 | P1 |  | SR | slw[.] | Shift Left Word |
| 011111 $\ldots . . . . . . . . . . . .11100011010$. | X | 1 | 109 | PPC |  | SR | srad[.] | Shift Right Algebraic Dword |
|  | XS | I | 109 | PPC |  | SR | sradi[.] | Shift Right Algebraic Dword Immediate |
| 011111 $\ldots \ldots \ldots \ldots \ldots .1100011000$. | X | I | 107 | P1 |  | SR | sraw[.] | Shift Right Algebraic Word |
| 011111 $\ldots \ldots . . . . . . . . . .1100111000$. | X | I | 107 | P1 |  | SR | srawi[.] | Shift Right Algebraic Word Immediate |
| 011111 $\ldots . . . . . . . . . . .1000011011$. | X | I | 108 | PPC |  | SR | srd[.] | Shift Right Dword |
| 011111 $\ldots \ldots . \ldots \ldots . . . . . .1000011000$. | X | I | 106 | P1 |  | SR | Stw[.] | Shift Right Word |
|  | D | 1 | 55 | P1 |  |  | stb | Store Byte |
|  | X | 111 | 967 | v2.05 | H |  | stbcix | Store Byte Caching Inhibited Indexed |
| 011111..... $\ldots$... $\ldots$.... 10101101101 | X | 11 | 870 | v2.06 |  |  | stbcx. | Store Byte Conditional Indexed \& record |
| 100111 ......... .............. | D | I | 55 | P1 |  |  | stbu | Store Byte with Update |
|  | X | I | 55 | P1 |  |  | stbux | Store Byte with Update Indexed |
| 011111.............. 00110101111 | X | I | 55 | P1 |  |  | stbx | Store Byte Indexed |
|  | DS | I | 58 | PPC |  |  | std | Store Dword |
| 011111.............. 10111001101 | X | 11 | 866 | v3.0 |  |  | stdat | Store Dword ATomic |
| 011111.............. 10100101001 | X | I | 62 | v2.06 |  |  | stdbrx | Store Dword Byte-Reverse Indexed |
| 011111..... $\ldots$... $\ldots$.... 11111101011 | X | III | 967 | v2.05 | H |  | stdcix | Store Dword Caching Inhibited Indexed |
| 011111 $\ldots \ldots . \ldots . . . . . . .00110101101$ | X | 11 | 873 | PPC |  |  | stdcx. | Store Dword Conditional Indexed \& record |
|  | DS | I | 58 | PPC |  |  | stdu | Store Dword with Update |
| 011111..... $\ldots$... $\ldots . . .00101101011$ | X | I | 58 | PPC |  |  | stdux | Store Dword with Update Indexed |
| 011111............. 00100101011 | X | I | 58 | PPC |  |  | stdx | Store Dword Indexed |
|  | D | 1 | 147 | P1 |  |  | stfd | Store Floating Double |
|  | DS | 1 | 150 | v2.05 |  |  | stfdp | Store Floating Double Pair |
| $011111 \ldots . . . . . . . . . . . .11100101111$ | X | 1 | 150 | v2.05 |  |  | stfdpx | Store Floating Double Pair Indexed |

Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 7 of 17)

| Instruction ${ }^{\mathbf{1}}$ <br> $012345 \quad 67890111111111222222222233$ <br> 123489012345678901 |  | $\begin{aligned} & \text { ᄃ } \\ & \text { ó } \\ & \text { in } \end{aligned}$ | $\begin{aligned} & \mathbb{\otimes} \\ & \stackrel{\pi}{\circ} \end{aligned}$ |  |  | $\begin{aligned} & r_{0}^{0} \\ & 0 \\ & 0 \\ & 0 \\ & 0 \\ & 0 \\ & \mathbf{0} \\ & \dot{\Sigma} \end{aligned}$ |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 110111 ..... ..... ..... ..... | D | 1 | 147 | P1 |  |  | sttdu | Store Floating Double with Update |
| 011111 ..... ..... ..... 10111101111 | X | 1 | 147 | P1 |  |  | stfdux | Store Floating Double with Update Indexed |
| $011111 \ldots \ldots . . . .10 . .10110101111$ | X | 1 | 147 | P1 |  |  | stfdx | Store Floating Double Indexed |
|  | X | 1 | 148 | PPC |  |  | stfiwx | Store Floating as Integer Word Indexed |
| $110100 \ldots \ldots . . . . . . . . . . . . . . . . .$. | D | 1 | 146 | P1 |  |  | stfs | Store Floating Single |
| 110101 ............... ..... ...... | D | I | 146 | P1 |  |  | stfsu | Store Floating Single with Update |
|  | X | 1 | 146 | P1 |  |  | stfsux | Store Floating Single with Update Indexed |
| 011111 .......... ..... 10100101111 | X | 1 | 146 | P1 |  |  | stfsx | Store Floating Single Indexed |
| 101100 .................. | D | I | 56 | P1 |  |  | sth | Store Hword |
|  | X | 1 | 61 | P1 |  |  | sthbrx | Store Hword Byte-Reverse Indexed |
| 011111 .......... ..... 11101101011 | X | 111 | 967 | v2.05 | H |  | sthcix | Store Hword Caching Inhibited Indexed |
| $011111 \ldots \ldots \ldots \ldots . .10110101101$ | X | 11 | 871 | v2.06 |  |  | sthcx. | Store Hword Conditional Indexed \& record |
| 101101 ..... ..... ..... ..... ..... | D | I | 56 | P1 |  |  | sthu | Store Hword with Update |
|  | X | 1 | 56 | P1 |  |  | sthux | Store Hword with Update Indexed |
|  | X | 1 | 56 | P1 |  |  | sthx | Store Hword Indexed |
| 101111 ..... ..... ..... ..... ..... | D | 1 | 63 | P1 |  |  | stmw | Store Multiple Word |
| 010011 \||||| ||1|| ||||| 01011100101 | XL | 111 | 957 | v3.0 | P |  | stop | Stop |
| 111110 .......... ..... ..... .... 10 | DS | 1 | 60 | v2.03 |  |  | stq | Store Qword |
| $011111 \ldots . . . . . . . . . . .00101101101$ | X | 11 | 876 | v2.07 |  |  | stacx. | Store Qword Conditional Indexed \& record |
| 011111 ..... .......... 10110101011 | X | 1 | 66 | P1 |  |  | stswi | Store String Word Immediate |
|  | X | 1 | 66 | P1 |  |  | stswx | Store String Word Indexed |
|  | X | 1 | 247 | v2.03 |  |  | stvebx | Store Vector Element Byte Indexed |
| $011111 \ldots . . . . . . . . .00101001111$ | X | 1 | 247 | v2.03 |  |  | stvehx | Store Vector Element Hword Indexed |
| $011111 \ldots . . . . . . . . . .00110001111$ | X | 1 | 248 | v2.03 |  |  | stvewx | Store Vector Element Word Indexed |
| $011111 \ldots . . . . . . . . .00111001111$ | X | 1 | 248 | v2.03 |  |  | stvx | Store Vector Indexed |
| $011111 \ldots . . . . . . . . . .01111001111$ | X | 1 | 248 | v2.03 |  |  | stux\| | Store Vector Indexed Last |
| 100100 ..... ..... ... | D | I | 57 | P1 |  |  | stw | Store Word |
| $011111 \ldots \ldots \ldots \ldots . . .10110001101$ | X | 11 | 866 | v3.0 |  |  | stwat | Store Word ATomic |
| $011111 \ldots \ldots \ldots . . . . .10100101101$ | X | 1 | 61 | P1 |  |  | stwbrx | Store Word Byte-Reverse Indexed |
| $011111 \ldots . . . . . . . . . .11100101011$ | X | 111 | 967 | v2.05 | H |  | stwcix | Store Word Caching Inhibited Indexed |
| $011111 \ldots . . . . . . . . . . .00100101101$ | X | 11 | 872 | PPC |  |  | stwcx. | Store Word Conditional Indexed \& record |
|  | D | 1 | 57 | P1 |  |  | stwu | Store Word with Update |
| $011111 \ldots . . . . . . . . . .00101101111$ | X | I | 57 | P1 |  |  | stwux | Store Word with Update Indexed |
| 011111 ..... ..... .... 00100 101111 | X | 1 | 57 | P1 |  |  | stwx | Store Word Indexed |
| 111101 ..... ..... ..... ..... .... 10 | DS |  | 499 | v3.0 |  |  | stxsd | Store VSX Scalar Dword |
| $011111 \ldots \ldots . . . . . . .1011001100$. | XX1 | 1 | 499 | v2.06 |  |  | stxsdx | Store VSX Scalar Dword Indexed |
|  | XX1 | 1 | 500 | v3.0 |  |  | stxsibx | Store VSX Scalar as Integer Byte Indexed |
| $011111 \ldots \ldots . . . . . . . .1110101101$. | XX1 | 1 | 500 | v3.0 |  |  | stxsihx | Store VSX Scalar as Integer Hword Indexed |
| $011111 \ldots \ldots \ldots \ldots . . . . .0010001100$. | XX1 | 1 | 501 | v2.07 |  |  | stxsiwx | Store VSX Scalar as Integer Word Indexed |
| 111101 ..... ..... ..... ..... ...11 | DS | I | 502 | v3.0 |  |  | stxssp | Store VSX Scalar SP |
| $011111 \ldots \ldots \ldots \ldots . . . . .1010001100$. | XX1 |  | 503 | v2.07 |  |  | stxsspx | Store VSX Scalar SP Indexed |
| 111101 ..... ..... ..... ..... ... 101 | DQ | I | 508 | v3.0 |  |  | sttv | Store VSX Vector |
|  | XX1 | 1 | 504 | v3.0 |  |  | stxvb16x | Store VSX Vector Byte*16 Indexed |
|  | XX1 | 1 | 505 | v2.06 |  |  | stxvd2x | Store VSX Vector Dword*2 Indexed |
|  | XX1 | 1 | 506 | v3.0 |  |  | stxvh8x | Store VSX Vector Hword*8 Indexed |
|  | XX1 | 1 | 508 | v3.0 |  |  | stxvl | Store VSX Vector with Length |
| $011111 \ldots \ldots . . . . . . . .0110101101$. | XX1 | 1 | 510 | v3.0 |  |  | stxvill | Store VSX Vector Left-justified with Length |
| $011111 \ldots \ldots \ldots \ldots . . .1110001100$. | XX1 | I | 507 | v2.06 |  |  | stxww4x | Store VSX Vector Word*4 Indexed |
|  | XX1 | I | 511 | v3.0 |  |  | stxvx | Store VSX Vector Indexed |
|  | XO | 1 | 70 | PPC |  | SR | subb[0][.] | Subtract From |
|  | XO | 1 | 71 | P1 |  | SR | subf[0][].] | Subtract From Carrying |
| 011111 ..... .......... 010001000. | XO | I | 72 | P1 |  | SR | subfe[0][].] | Subtract From Extended |
| $001000 \ldots \ldots . .$. | D | 1 | 71 | P1 |  | SR | subfic | Subtract From Immediate Carrying |
| $011111 \ldots . . . . . .11111 .011101000$. | XO | 1 | 72 | P1 |  | SR | subfme[0][.] | Subtract From Minus One Extended |
| $011111 \ldots \ldots . . .11111 .011001000$. | XO | 1 | 73 | P1 |  | SR | subfze[0][.] | Subtract From Zero Extended |
| 011111 \|11.. |1111 ||||| 10010101101 | X | 11 | 877 | P1 |  |  | sync | Synchronize |

Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 8 of 17)


Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 9 of 17)

|  Instruction  <br> 1   <br> 1 11111 11112 <br> 01234567890 12345 67890 | $\begin{aligned} & 22222 \\ & 12345 \end{aligned}$ | $\begin{aligned} & 222233 \\ & 678901 \end{aligned}$ |  | $\begin{aligned} & \text { ̌i } \\ & \text { D } \end{aligned}$ | $\begin{aligned} & \mathbb{\otimes} \\ & \stackrel{\pi}{\circ} \end{aligned}$ |  |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $000100 \ldots . . . . . . . . . .$. | 1111 | 000110 | VC | 1 | 331 | v2.03 |  |  | vcmpbfp[.] | Vector Compare Bounds Floating-Point |
| $000100 \ldots . . . . . . . . . .$. | . 0011 | 000110 | VC | 1 | 332 | v2.03 |  |  | vcmpeqfp[] | Vector Compare Equal To Floating-Point |
| $000100 \ldots \ldots \ldots \ldots$ | . 0000 | 000110 | VC | 1 | 306 | v2.03 |  |  | vcmpequb[.] | Vector Compare Equal Unsigned Byte |
| $000100 \ldots \ldots . . . . . . . .$. | . 0011 | 000111 | VC | I | 307 | v2.07 |  |  | vcmpequd[.] | Vector Compare Equal Unsigned Dword |
| $000100 \ldots . . . . . . . . . .$. | . 0001 | 000110 | VC | 1 | 306 | v2.03 |  |  | vcmpequh[.] | Vector Compare Equal Unsigned Hword |
| $000100 \ldots . . . . . . . . . .$. | . 0010 | 000110 | VC | 1 | 307 | v2.03 |  |  | vcmpequw[]] | Vector Compare Equal Unsigned Word |
| $000100 \ldots . . . . . . . . .$. | . 0111 | 000110 | VC | I | 332 | v2.03 |  |  | vcmpgefp[] | Vector Compare Greater Than or Equal To Floating-Point |
| $000100 \ldots . . . . . . . . . .$. | 1011 | 000110 | VC | 1 | 333 | v2.03 |  |  | vcmpgtfp[] | Vector Compare Greater Than Floating-Point |
| 000100 ............... | 1100 | 000110 | VC | I | 308 | v2.03 |  |  | vcmpgtsb[.] | Vector Compare Greater Than Signed Byte |
| $000100 \ldots . . . . . . . . .$. | . 1111 | 000111 | VC | I | 308 | v2.07 |  |  | vcmpgtsd[]] | Vector Compare Greater Than Signed Dword |
| $000100 \ldots . . . . . . . . . .$. | 1101 | 000110 | VC | 1 | 309 | v2.03 |  |  | vcmpgtsh[.] | Vector Compare Greater Than Signed Hword |
| $000100 \ldots \ldots \ldots \ldots$ | 1110 | 000110 | VC | I | 309 | v2.03 |  |  | vcmpgtsw[.] | Vector Compare Greater Than Signed Word |
| 000100 ............... | . 1000 | 000110 | VC | I | 310 | v2.03 |  |  | vcmpgtub[.] | Vector Compare Greater Than Unsigned Byte |
| $000100 \ldots . . . . . . . . .$. | 1011 | 000111 | VC | 1 | 310 | v2.07 |  |  | vcmpgtud[]] | Vector Compare Greater Than Unsigned Dword |
| $000100 \ldots \ldots . . . . . .$. | 1001 | 000110 | VC | 1 | 311 | v2.03 |  |  | vcmpgtuh[.] | Vector Compare Greater Than Unsigned Hword |
| 000100 ............... | . 1010 | 000110 | VC | I | 311 | v2.03 |  |  | vcmpgtuw[.] | Vector Compare Greater Than Unsigned Word |
| $000100 \ldots . . . . . . . . .$. | . 0000 | 000111 | VC | 1 | 312 | v3.0 |  |  | vcmpneb[] | Vector Compare Not Equal Byte |
| $000100 \ldots \ldots \ldots . . . . . . .$. | . 0001 | 000111 | VC | 1 | 313 | v3.0 |  |  | vcmpneh[.] | Vector Compare Not Equal Hword |
| $000100 \ldots \ldots \ldots \ldots$ | . 0010 | 000111 | VC | I | 314 | v3.0 |  |  | vcmpnew[] | Vector Compare Not Equal Word |
| $000100 \ldots . . . . . . . . . .$. | . 0100 | 000111 | VC | 1 | 312 | v3.0 |  |  | vcmpnezb[.] | Vector Compare Not Equal or Zero Byte |
| $000100 \ldots \ldots \ldots \ldots$ | . 0101 | 000111 | VC | 1 | 313 | v3.0 |  |  | vcmpnezh[.] | Vector Compare Not Equal or Zero Hword |
| $000100 \ldots \ldots \ldots$ | . 0110 | 000111 | VC | 1 | 314 | v3.0 |  |  | vcmpnezw[] | Vector Compare Not Equal or Zero Word |
| $000100 \ldots . . . . . . . . . .$. | 01111 | 001010 | VX | I | 327 | v2.03 |  |  | vctsxs | Vector Convert To Signed Word Saturate |
| $000100 \ldots . . . . . . . . .$. | 01110 | 001010 | VX | 1 | 327 | v2.03 |  |  | vctuxs | Vector Convert To Unsigned Word Saturate |
| $000100 \ldots . .11100 \ldots \ldots$ | 11000 | 000010 | VX | 1 | 344 | v3.0 |  |  | vctzb | Vector Count Trailing Zeros Byte |
| $000100 \ldots . .11111 \ldots$ | 11000 | 000010 | VX | I | 344 | v3.0 |  |  | vctzd | Vector Count Trailing Zeros Dword |
| $000100 \ldots . .11101 \ldots .$. | 11000 | 000010 | VX | I | 344 | v3.0 |  |  | vctzh | Vector Count Trailing Zeros Hword |
| $000100 \ldots . .00001 \ldots .$. | 11000 | 000010 | VX | 1 | 345 | v3.0 |  |  | vctzlsbb | Vector Count Trailing Zero Least-Significant Bits Byte |
| $000100 \ldots . .11110 \ldots .$. | 11000 | 000010 | VX | + | 344 | v3.0 |  |  | vctzw | Vector Count Trailing Zeros Word |
| $000100 \ldots . . . . . . . . . . . . .$. | 11010 | 000100 | VX | 1 | 315 | v2.07 |  |  | veqv | Vector Logical Equivalence |
| $000100 \ldots . .11111 . .$. | 00110 | 001010 | VX | 1 | 334 | v2.03 |  |  | vexptefp | Vector 2 Raised to the Exponent Estimate Floating-Point |
| $000100 \ldots . .1 \ldots . . . .$. | 01011 | 001101 | VX | 1 | 269 | v3.0 |  |  | vextractd | Vector Extract Dword |
| $000100 \ldots . .1 \ldots . . . . . .$. | 01000 | 001101 | VX | 1 | 269 | v3.0 |  |  | vextractub | Vector Extract Unsigned Byte |
| $000100 \ldots . .1 . . . . . . .$. | 01001 | 001101 | VX | 1 | 269 | v3.0 |  |  | vextractuh | Vector Extract Unsigned Hword |
| $000100 \ldots . .1 \ldots . . .$. | 01010 | 001101 | VX | 1 | 269 | v3.0 |  |  | vextractuw | Vector Extract Unsigned Word |
| $000100 \ldots \ldots .11000 \ldots .$. | 11000 | 000010 | VX | 1 | 296 | v3.0 |  |  | vextsb2d | Vector Extend Sign Byte to Dword |
| $000100 \ldots . .10000 \ldots .$. | 11000 | 000010 | VX | I | 296 | v3.0 |  |  | vextsb2w | Vector Extend Sign Byte to Word |
| $000100 \ldots . .11001 \ldots .$. | 11000 | 000010 | VX | 1 | 296 | v3.0 |  |  | vextsh2d | Vector Extend Sign Hword to Dword |
| $000100 \ldots . . .10001 \ldots .$. | 11000 | 000010 | VX | 1 | 296 | v3.0 |  |  | vextsh2w | Vector Extend Sign Hword to Word |
| $000100 \ldots . .11010 \ldots .$. | 11000 | 000010 | VX | I | 297 | v3.0 |  |  | vextsw2d | Vector Extend Sign Word to Dword |
| $000100 \ldots . . . . . . . . . .$. | 11000 | 001101 | VX | + | 346 | v3.0 |  |  | vextublx | Vector Extract Unsigned Byte Left-Indexed |
| $000100 \ldots . . . . . . . . . . .$. | 11100 | 001101 | VX | 1 | 346 | v3.0 |  |  | vextubrx | Vector Extract Unsigned Byte Right-Indexed |
| $000100 \ldots . . . . . . . . .$. | 11001 | 001101 | VX | 1 | 346 | v3.0 |  |  | vextuhlx | Vector Extract Unsigned Hword Left-Indexed |
| $000100 \ldots \ldots . . . . . . .$. | 11101 | 001101 | VX | + | 346 | v3.0 |  |  | vextuhrx | Vector Extract Unsigned Hword Right-Indexed |
| $000100 \ldots . . . . . . . . . . . .$. | 11010 | 001101 | VX | 1 | 347 | v3.0 |  |  | vextuwlx | Vector Extract Unsigned Word Left-Indexed |
| $000100 \ldots . . . . . . . . . .$. | 11110 | 001101 | VX | 1 | 347 | v3.0 |  |  | vextuwrx | Vector Extract Unsigned Word Right-Indexed |
| $000100 \ldots . .11111 \ldots$ | 10100 | 001100 | VX | 1 | 342 | v2.07 |  |  | vgbbd | Vector Gather Bits by Byte by Dword |
| $000100 \ldots . .1 . . . . . . .$. | 01100 | 001101 | VX | , | 270 | v3.0 |  |  | vinsertb | Vector Insert Byte |
| $000100 \ldots . .1 . . . . . .$. | 01111 | 001101 | VX | I | 270 | v3.0 |  |  | vinsertd | Vector Insert Dword |
| $000100 \ldots . . .1 . . . . . . .$. | 01101 | 001101 | VX | , | 270 | v3.0 |  |  | vinserth | Vector Insert Hword |
| $000100 \ldots \ldots 1 \ldots \ldots$ | 01110 | 001101 | VX | 1 | 270 | v3.0 |  |  | vinsertw | Vector Insert Word |
| $000100 \ldots . . . . . . . . .$. | 00111 | 001010 | VX | I | 334 | v2.03 |  |  | vogefp | Vector Log Base 2 Estimate Floating-Point |
| $000100 \ldots . . . . . . . .$. | ..... | 101110 | VA | 1 | 325 | v2.03 |  |  | vmaddfp | Vector Multiply-Add Floating-Point |
| $000100 \ldots . . . . . . . . . . .$. | 10000 | 001010 | VX | 1 | 326 | v2.03 |  |  | vmaxfp | Vector Maximum Floating-Point |
| $000100 \ldots \ldots \ldots . . . . . . .$. | 00100 | 000010 | VX | 1 | 302 | v2.03 |  |  | vmaxsb | Vector Maximum Signed Byte |
| 000100 ..... ........... | 00111 | 000010 | VX | I | 302 | v2.07 |  |  | vmaxsd | Vector Maximum Signed Dword |

Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 10 of 17)


Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 11 of 17)

## I



Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 12 of 17)
I


Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 13 of 17)


Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 14 of 17)

| Instruction $^{\mathbf{1}}$ <br> $01234567890 \quad 1111111112 \quad 22222 \quad 222233$ <br> $1234567890 \quad 12345678901$ |  | $\begin{aligned} & \text { ro } \\ & \text { on } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & \underset{\sim}{0} \end{aligned}$ |  |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 111111............. 0000100100. | X | 1 | 604 | v3.0 |  |  | xsmulqp[0] | VSX Scalar Multiply QP |
| 111100 ............. $00010000 \ldots$ | XX3 | I | 606 | v2.07 |  |  | xsmulsp | VSX Scalar Multiply SP |
| $111100 \ldots . .11111 \ldots . .10110$ 1001.. | XX2 | I | 608 | v2.06 |  |  | xsnabsdp | VSX Scalar Negative Absolute DP |
| 111111.... $01000 \ldots . .11001001001$ | X | I | 608 | v3.0 |  |  | xsnabsqp | VSX Scalar Negative Absolute QP |
| 111100 .... 11111 .... 10111 1001.. | XX2 | I | 609 | v2.06 |  |  | xsnegdp | VSX Scalar Negate DP |
| 111111.... $10000 \ldots . .11001001001$ | X | I | 609 | v3.0 |  |  | xsnegqp | VSX Scalar Negate QP |
| $111100 \ldots \ldots . . . . . . . .10100001 \ldots$ | XX3 | 1 | 610 | v2.06 |  |  | xsnmaddadp | VSX Scalar Negative Multiply-Add Type-A DP |
| 111100 $\ldots . . . \ldots \ldots . . . . .10000001 \ldots$ | XX3 | I | 615 | v2.07 |  |  | xsnmaddasp | VSX Scalar Negative Multiply-Add Type-A SP |
| 111100 $\ldots . . . . . . . . . . .10101001 \ldots$ | XX3 | I | 610 | v2.06 |  |  | xsnmaddmdp | VSX Scalar Negative Multiply-Add Type-M DP |
| 111100 $\ldots \ldots . \ldots \ldots . . . . . .10001001 \ldots$ | XX3 | I | 615 | v2.07 |  |  | xsnmaddmsp | VSX Scalar Negative Multiply-Add Type-M SP |
| 111111 $\ldots .$. $\ldots .$. $\ldots .$. 01110 00100 | X | I | 618 | v3.0 |  |  | xsnmaddqp[0] | VSX Scalar Negative Multiply-Add QP |
|  | XX3 | I | 621 | v2.06 |  |  | xsnmsubadp | VSX Scalar Negative Multiply-Subtract Type-A DP |
| 111100 .............. 10010 001... | XX3 | 1 | 624 | v2.07 |  |  | xsnmsubasp | VSX Scalar Negative Multiply-Subtract Type-A SP |
| 111100 $\ldots . . . \ldots . . . . . .10111001 \ldots$ | XX3 | I | 621 | v2.06 |  |  | xsnmsubmdp | VSX Scalar Negative Multiply-Subtract Type-M DP |
| 111100 $\ldots . . . \ldots . . . . . . .10011001 \ldots$ | XX3 | I | 624 | v2.07 |  |  | xsnmsubmsp | VSX Scalar Negative Multiply-Subtract Type-M SP |
| 111111 ..... .... ..... 0111100100. | X | I | 627 | v3.0 |  |  | xsnmsubqp[0] | VSX Scalar Negative Multiply-Subtract QP |
| $111100 \ldots . .11111 \ldots . .001001001 .$. | XX2 | I | 630 | v2.06 |  |  | xsrdpi | VSX Scalar Round DP to Integral to Nearest Away |
| 111100 $\ldots . .11111 \ldots . .00110$ 1011.. | XX2 | I | 631 | v2.06 |  |  | xsrdpic | VSX Scalar Round DP to Integral using Current rounding mode |
| 111100 $\ldots . . .11111 \ldots . . .001111001 .$. | XX2 | I | 632 | v2.06 |  |  | xsrdpim | VSX Scalar Round DP to Integral toward -Infinity |
| $111100 \ldots . .11111 \ldots . .00110$ 1001.. | XX2 | 1 | 632 | v2.06 |  |  | xsrdpip | VSX Scalar Round DP to Integral toward + Infinity |
| 111100 $\ldots . .11111 \ldots . .00101$ 1001.. | XX2 | I | 633 | v2.06 |  |  | xsrdpiz | VSX Scalar Round DP to Integral toward Zero |
| 111100 $\ldots . .11111 \ldots . . .001011010 .$. | XX2 | I | 634 | v2.06 |  |  | xsredp | VSX Scalar Reciprocal Estimate DP |
| 111100 $\ldots . .11111 \ldots . . .000011010 .$. | XX2 | I | 635 | v2.07 |  |  | xsresp | VSX Scalar Reciprocal Estimate SP |
| 111111.... 1111. ...... 00000101. | Z23 | 1 | 636 | v3.0 |  |  | xsrqpi[x] | VSX Scalar Round QP to Integral |
| 111111..... 1111. ..... . $00100101 /$ | Z23 | I | 638 | v3.0 |  |  | xsrqpxp | VSX Scalar Round QP to XP |
| 111100 $\ldots . .11111 \ldots . . .100011001 .$. | XX2 | I | 640 | v2.07 |  |  | xsrsp | VSX Scalar Round DP to SP |
| $111100 \ldots . .11111 \ldots . .00100$ 1010.. | XX2 | 1 | 641 | v2.06 |  |  | xsrsqrtedp | VSX Scalar Reciprocal Square Root Estimate DP |
| $111100 \ldots . .11111 \ldots . .00000$ 1010.. | XX2 | I | 642 | v2.07 |  |  | xsrsqrtesp | VSX Scalar Reciprocal Square Root Estimate SP |
| 111100 $\ldots . .11111 \ldots . .1001001011 .$. | XX2 | I | 643 | v2.06 |  |  | xssqritdp | VSX Scalar Square Root DP |
| 111111 $\ldots . . .11011 \ldots . . .1100100100$. | X | I | 644 | v3.0 |  |  | xssqrtqp[0] | VSX Scalar Square Root QP |
| 111100 $\ldots . .11111 \ldots . . .00000$ 1011.. | XX2 | I | 646 | v2.07 |  |  | xssqrtsp | VSX Scalar Square Root SP |
| 111100 $\ldots . . . \ldots . . . . . . .00101000 \ldots$ | XX3 | I | 647 | v2.06 |  |  | xssubdp | VSX Scalar Subtract DP |
| 111111............. 1000000100. | X | I | 649 | v3.0 |  |  | xssubqp[0] | VSX Scalar Subtract QP |
| 111100 $\ldots . . . . . . . . . . . . ~ 00001000 \ldots$ | XX3 | 1 | 651 | v2.07 |  |  | xssubsp | VSX Scalar Subtract SP |
| $111100 \ldots 11 \ldots . . . . . .00111101 . .1$ | XX3 | I | 653 | v2.06 |  |  | xstdivdp | VSX Scalar Test for software Divide DP |
| $111100 \ldots 1111111 \ldots \ldots 001101010.1$ | XX2 | I | 654 | v2.06 |  |  | xstsqrtdp | VSX Scalar Test for software Square Root DP |
|  | XX2 | I | 655 | v3.0 |  |  | xststdcdp | VSX Scalar Test Data Class DP |
| 111111 $\ldots . . . . . . . . . . . . .10110001001$ | X | I | 656 | v3.0 |  |  | xststdcqp | VSX Scalar Test Data Class QP |
| 111100 $\ldots \ldots . \ldots \ldots . . . . . .100101010 .1$ | XX2 | I | 657 | v3.0 |  |  | xststdcsp | VSX Scalar Test Data Class SP |
| 111100 $\ldots . . .00000 \ldots . .101011011 .1$ | XX2 | I | 658 | v3.0 |  |  | xsxexpdp | VSX Scalar Extract Exponent DP |
| 111111 $\ldots . . .00010 \ldots . .11001001001$ | X | 1 | 658 | v3.0 |  |  | xsxexpqp | VSX Scalar Extract Exponent QP |
| 111100 $\ldots . . .00001 \ldots . .101011011 .1$ | XX2 | I | 659 | v3.0 |  |  | xsxsigdp | VSX Scalar Extract Significand DP |
| 111111 $\ldots . . .10010 \ldots . .11001001001$ | X | 1 | 659 | v3.0 |  |  | xsxsigqp | VSX Scalar Extract Significand QP |
| 111100 $\ldots . . .11111 \ldots . . .111011001 .$. | XX2 | 1 | 660 | v2.06 |  |  | xvabsdp | VSX Vector Absolute DP |
| 111100 $\ldots . .11111 \ldots . .1110011001 .$. | XX2 | I | 660 | v2.06 |  |  | xvabssp | VSX Vector Absolute SP |
| 111100 $\ldots . . . \ldots . . . . . . .01100000 \ldots$ | XX3 | 1 | 661 | v2.06 |  |  | xvadddp | VSX Vector Add DP |
|  | XX3 | I | 665 | v2.06 |  |  | xvaddsp | VSX Vector Add SP |
|  | XX3 | I | 667 | v2.06 |  |  | xvcmpeqdp[.] | VSX Vector Compare Equal DP |
|  | XX3 | I | 668 | v2.06 |  |  | xvcmpeqsp[.] | VSX Vector Compare Equal SP |
|  | XX3 | I | 669 | v2.06 |  |  | xvcmpgedp[.] | VSX Vector Compare Greater Than or Equal DP |
|  | XX3 | 1 | 670 | v2.06 |  |  | xvcmpgesp[.] | VSX Vector Compare Greater Than or Equal SP |
| 111100 $\ldots . . . . . . . . . . . .1101011 . .$. | XX3 | 1 | 671 | v2.06 |  |  | xvcmpgtdp[.] | VSX Vector Compare Greater Than DP |
| 111100 $\ldots . . . \ldots . . . . . . . .1001011 \ldots$ | XX3 | 1 | 672 | v2.06 |  |  | xvcmpgtsp[.] | VSX Vector Compare Greater Than SP |
| 111100 $\ldots . . . . . . . . . . . . . .11111011 . .$. | XX3 | I | 673 | v3.0 |  |  | xvcmpnedp[.] | VSX Vector Compare Not Equal Double-Precision |
| 111100 $\ldots \ldots . . . . . . . . . . . .10111011 . .$. | XX3 | I | 674 | v3.0 |  |  | xvcmpnesp[.] | VSX Vector Compare Not Equal Single-Precision |
| 111100 ............... 11110000... | XX3 | I | 675 | v2.06 |  |  | xvcpsgndp | VSX Vector Copy Sign DP |

Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 15 of 17)

| Instruction ${ }^{\mathbf{1}}$ |  |  |  |  | $\begin{aligned} & \text { पे } \\ & \text { © } \end{aligned}$ | $\begin{aligned} & \mathbb{Z} \\ & \stackrel{\pi}{\circ} \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \text { 흐N } \\ & \stackrel{n}{0} \end{aligned}$ |  | $\begin{aligned} & \mathrm{C}_{0}^{0} \\ & 0 \\ & 0 \\ & 0 \\ & 0 \\ & \dot{\Sigma} \end{aligned}$ |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 012345 | 678901234567890 | $\begin{aligned} & 22222 \\ & 12345 \end{aligned}$ | $\begin{aligned} & 222233 \\ & 678901 \end{aligned}$ |  |  |  |  |  |  |  |  |
| 111100 | - | 11010 |  | XX3 | 1 | 675 | v2.06 |  |  | xvcpsgnsp | VSX Vector Copy Sign SP |
| 111100 | $\ldots . .11111 . .$. | 11000 |  | XX2 | I | 676 | v2.06 |  |  | xvcvdpsp | VSX Vector Convert DP to SP |
| 111100 | $\ldots . . .11111 . .$. | 11101 | 1000 | XX2 | I | 677 | v2.06 |  |  | xvcvdpsxds | VSX Vector Convert DP to Signed Dword truncate |
| 111100 | ..... 11111..... | 01101 |  | XX2 | I | 679 | v2.06 |  |  | xvcvdpsxws | VSX Vector Convert DP to Signed Word truncate |
| 111100 | ..... \|111| ..... | 11100 |  | XX2 | I | 681 | v2.06 |  |  | xvcvdpuxds | VSX Vector Convert DP to Unsigned Dword truncate |
| 111100 | ..... 11111..... | 01100 |  | XX2 | I | 683 | v2.06 |  |  | xvcvdpuxws | VSX Vector Convert DP to Unsigned Word truncate |
| 111100 | $\ldots . .11000 \ldots$. | 11101 |  | XX2 | I | 685 | v3.0 |  |  | xvcvhpsp | VSX Vector Convert HP to SP |
| 111100 | ..... \|1111 ..... | 11100 |  | XX2 | I | 686 | v2.06 |  |  | xvcvspdp | VSX Vector Convert SP to DP |
| 111100 | ..... $11001 \ldots .$. | 11101 |  | XX2 | I | 687 | v3.0 |  |  | xvcvsphp | VSX Vector Convert SP to HP |
| 111100 | ..... $\\| 1111 \ldots$ | 11001 |  | XX2 | 1 | 688 | v2.06 |  |  | xvcvspsxds | VSX Vector Convert SP to Signed Dword truncate |
| 111100 | ..... 11111.... | 01001 | 1000 | XX2 | I | 690 | v2.06 |  |  | xvcvspsxws | VSX Vector Convert SP to Signed Word truncate |
| 111100 | $\ldots$....\|1111.... | 11000 |  | XX2 | I | 692 | v2.06 |  |  | xvcvspuxds | VSX Vector Convert SP to Unsigned Dword truncate |
| 111100 | ..... \|l|11 ..... | 01000 |  | XX2 | I | 694 | v2.06 |  |  | xvcvspuxws | VSX Vector Convert SP to Unsigned Word truncate |
| 111100 | ..... \|1111..... | 11111 |  | XX2 | I | 696 | v2.06 |  |  | xvcvsxddp | VSX Vector Convert Signed Dword to DP |
| 111100 | ..... \|1111..... | 11011 |  | XX2 | I | 696 | v2.06 |  |  | xvcvsxdsp | VSX Vector Convert Signed Dword to SP |
| 111100 | ..... \|l|11..... | 01111 |  | XX2 | I | 697 | v2.06 |  |  | xvCVsxwdp | VSX Vector Convert Signed Word to DP |
| 111100 | ..... \|1111..... | 01011 |  | XX2 | I | 697 | v2.06 |  |  | xvcvsxwsp | VSX Vector Convert Signed Word to SP |
| 111100 | $\ldots . .11111 . .$. | 11110 |  | XX2 | I | 698 | v2.06 |  |  | xvcvuxddp | VSX Vector Convert Unsigned Dword to DP |
| 111100 | ..... \|1111..... | 11010 |  | XX2 | I | 698 | v2.06 |  |  | xvcvuxdsp | VSX Vector Convert Unsigned Dword to SP |
| 111100 | ..... 11111..... | 01110 |  | XX2 | I | 699 | v2.06 |  |  | xvcvuxwdp | VSX Vector Convert Unsigned Word to DP |
| 111100 | .....11111.... | 01010 |  | XX2 | 1 | 699 | v2.06 |  |  | xvcvuxwsp | VSX Vector Convert Unsigned Word to SP |
| 111100 | ..... ........... | 011110 | 000... | XX3 | I | 700 | v2.06 |  |  | xvdivdp | VSX Vector Divide DP |
| 111100 | ............... | 01011 | 000... | XX3 | I | 702 | v2.06 |  |  | xvdivsp | VSX Vector Divide SP |
| 111100 | - | 111110 |  | XX3 | I | 704 | v3.0 |  |  | xviexpdp | VSX Vector Insert Exponent DP |
| 111100 | ............... | 11011 | 000... | XX3 | I | 704 | v3.0 |  |  | xviexpsp | VSX Vector Insert Exponent SP |
| 111100 | ..... .......... | 01100 | 001... | XX3 | I | 705 | v2.06 |  |  | xvmaddadp | VSX Vector Multiply-Add Type-A DP |
| 111100 | ........ | 01000 |  | XX3 | I | 708 | v2.06 |  |  | xvmaddasp | VSX Vector Multiply-Add Type-A SP |
| 111100 | .............. | 01101 | 001... | XX3 | I | 705 | v2.06 |  |  | xvmaddmdp | VSX Vector Multiply-Add Type-M DP |
| 111100 | ............... | 01001 | 001... | XX3 | I | 708 | v2.06 |  |  | xvmaddmsp | VSX Vector Multiply-Add Type-M SP |
| 111100 | . ......... | 11100 |  | XX3 | I | 711 | v2.06 |  |  | xvmaxdp | VSX Vector Maximum DP |
| 111100 | ..... ........... | 11000 | 000... | XX3 | I | 713 | v2.06 |  |  | xvmaxsp | VSX Vector Maximum SP |
| 111100 | ..... .......... | 11101 | 000... | XX3 | I | 715 | v2.06 |  |  | xvmindp | VSX Vector Minimum DP |
| 111100 | ........... | 11001 |  | XX3 | I | 717 | v2.06 |  |  | xvminsp | VSX Vector Minimum SP |
| 111100 | ................ | 01110 | 001... | XX3 | I | 719 | v2.06 |  |  | xvmsubadp | VSX Vector Multiply-Subtract Type-A DP |
| 111100 | ..... .......... | 01010 | 001... | XX3 | I | 722 | v2.06 |  |  | xvmsubasp | VSX Vector Multiply-Subtract Type-A SP |
| 111100 | .............. | 01111 | 001... | XX3 | 1 | 719 | v2.06 |  |  | xvmsubmdp | VSX Vector Multiply-Subtract Type-M DP |
| 111100 | .... .......... | 01011 | 001... | XX3 | I | 722 | v2.06 |  |  | xvmsubmsp | VSX Vector Multiply-Subtract Type-M SP |
| 111100 | $\ldots \ldots \ldots \ldots$ | 01110 |  | XX3 | I | 725 | v2.06 |  |  | xvmuldp | VSX Vector Multiply DP |
| 111100 | - | 01010 |  | XX3 | I | 727 | v2.06 |  |  | xvmulsp | VSX Vector Multiply SP |
| 111100 | ..... \|1111..... | 11110 |  | XX2 | I | 729 | v2.06 |  |  | xvnabsdp | VSX Vector Negative Absolute DP |
| 111100 | ..... \|1111 ..... | 11010 |  | XX2 | I | 729 | v2.06 |  |  | xvnabssp | VSX Vector Negative Absolute SP |
| 111100 | ..... 11111..... | 11111 |  | XX2 | I | 730 | v2.06 |  |  | xvnegdp | VSX Vector Negate DP |
| 111100 | ..... \|1111..... | 11011 |  | XX2 | I | 730 | v2.06 |  |  | xvnegsp | VSX Vector Negate SP |
| 111100 | $\ldots . . . . . . . .$. | 11100 |  | XX3 | I | 731 | v2.06 |  |  | xvnmaddadp | VSX Vector Negative Multiply-Add Type-A DP |
| 111100 | .... .......... | 11000 | 001... | XX3 | I | 736 | v2.06 |  |  | xvnmaddasp | VSX Vector Negative Multiply-Add Type-A SP |
| 111100 | ... .......... | 11101 |  | XX3 | I | 731 | v2.06 |  |  | xvnmaddmdp | VSX Vector Negative Multiply-Add Type-M DP |
| 111100 | $\ldots \ldots \ldots \ldots$ | 11001 |  | XX3 | I | 736 | v2.06 |  |  | xvnmaddmsp | VSX Vector Negative Multiply-Add Type-M SP |
| 111100 | $\ldots \ldots \ldots \ldots$ | 11110 | 001... | XX3 | I | 739 | v2.06 |  |  | xvnmsubadp | VSX Vector Negative Multiply-Subtract Type-A DP |
| 111100 | ..... .......... | 11010 | 001... | XX3 | I | 742 | v2.06 |  |  | xvnmsubasp | VSX Vector Negative Multiply-Subtract Type-A SP |
| 111100 | $\ldots \ldots \ldots$ | 11111 |  | XX3 | I | 739 | v2.06 |  |  | xvnmsubmdp | VSX Vector Negative Multiply-Subtract Type-M DP |
| 111100 | ............... | 11011 |  | XX3 | I | 742 | v2.06 |  |  | xvnmsubmsp | VSX Vector Negative Multiply-Subtract Type-M SP |
| 111100 | ..... 11111..... | 01100 | 1001. | XX2 | I | 745 | v2.06 |  |  | xvrdpi | VSX Vector Round DP to Integral to Nearest Away |
| 111100 | $\ldots . .11\\| \\| \ldots$ | 01110 | 1011. | XX2 | 1 | 745 | v2.06 |  |  | xvrdpic | VSX Vector Round DP to Integral using Current rounding mode |
| 111100 | ..... \|1111..... | 01111 | 1001. | XX2 | I | 746 | v2.06 |  |  | xvrdpim | VSX Vector Round DP to Integral toward -Infinity |
| 111100 | $\ldots . . .11111 \ldots$ | 01110 | 1001. | XX2 | 1 | 746 | v2.06 |  |  | xurdpip | VSX Vector Round DP to Integral toward +Infinity |
| 111100 | ..... \|1111..... | 01101 | 1001. | XX2 | I | 747 | v2.06 |  |  | xvrdpiz | VSX Vector Round DP to Integral toward Zero |

Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 16 of 17)

| Instruction ¹ $01234567890111111111222222 \quad 222233$ $1234567890 \quad 12345678901$ |  | $\begin{aligned} & \text { ro } \\ & \text { ón } \end{aligned}$ | $\begin{aligned} & \mathbb{0} \\ & \underset{\sim}{0} \end{aligned}$ | $\begin{aligned} & \text { N } \\ & \frac{0}{0} \\ & \text { N} \\ & \hline \end{aligned}$ |  |  |  | Name |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 111100 $\ldots . .11111 \ldots . . .011011010$. | XX2 | \| | 748 | v2.06 |  |  | xvredp | VSX Vector Reciprocal Estimate DP |
| $111100 \ldots . .11111 \ldots . . .01001$ 1010.. | XX2 | I | 749 | v2.06 |  |  | xvresp | VSX Vector Reciprocal Estimate SP |
| 111100 $\ldots . .11111 \ldots . . .1010001001 .$. | XX2 | I | 750 | v2.06 |  |  | xvrspi | VSX Vector Round SP to Integral to Nearest Away |
| $111100 \ldots . .11111 \ldots . .010101011$. | XX2 | , | 750 | v2.06 |  |  | xvrspic | VSX Vector Round SP to Integral using Current rounding mode |
| 111100 $\ldots . .11111 \ldots . . .010111001$. | XX2 | I | 751 | v2.06 |  |  | xvrspim | VSX Vector Round SP to Integral toward -Infinity |
| 111100 $\ldots . .11111 \ldots \ldots .1010101001 .$. | XX2 | I | 751 | v2.06 |  |  | xvrspip | VSX Vector Round SP to Integral toward +Infinity |
| $111100 \ldots . .11111 \ldots . .010011001 .$. | XX2 | I | 752 | v2.06 |  |  | xvrspiz | VSX Vector Round SP to Integral toward Zero |
| 111100 $\ldots . .11111 \ldots . . .1011001010$. | XX2 | 1 | 752 | v2.06 |  |  | xvrsqrtedp | VSX Vector Reciprocal Square Root Estimate DP |
| 111100 $\ldots . .11111 \ldots \ldots 010001010$. | XX2 | I | 754 | v2.06 |  |  | xvrsqriesp | VSX Vector Reciprocal Square Root Estimate SP |
| 111100 .... \|111| .... $011001011 .$. | XX2 | I | 755 | v2.06 |  |  | xvsqrtdp | VSX Vector Square Root DP |
| 111100 $\ldots . .11111 \ldots . .010001011$. | XX2 | I | 756 | v2.06 |  |  | xvsqrtsp | VSX Vector Square Root SP |
| 111100 $\ldots \ldots . . . . . . . . . . .01101000 \ldots$ | XX3 | 1 | 757 | v2.06 |  |  | xvsubdp | VSX Vector Subtract DP |
| 111100 $\ldots . . . . . . . . . . .01001000 \ldots$ | XX3 | 1 | 759 | v2.06 |  |  | xvsubsp | VSX Vector Subtract SP |
| 111100 $\ldots 111 \ldots . . . . . . .01111101 . .1$ | XX3 | I | 761 | v2.06 |  |  | xvtdivdp | VSX Vector Test for software Divide DP |
|  | XX3 | I | 762 | v2.06 |  |  | xvtdivsp | VSX Vector Test for software Divide SP |
| $111100 \ldots 1111111 \ldots . .011101010 .1$ | XX2 | I | 763 | v2.06 |  |  | xvtsqrtdp | VSX Vector Test for software Square Root DP |
| $111100 \ldots 1111111 \ldots \ldots .010101010 .1$ | XX2 | 1 | 763 | v2.06 |  |  | xvtsqrtsp | VSX Vector Test for software Square Root SP |
| 111100 .............. 1111. 101... | XX2 | I | 764 | v3.0 |  |  | xvtstdcdp | VSX Vector Test Data Class DP |
| 111100 $\ldots . . . . . . . . . . . .1101 .101 \ldots$ | XX2 | I | 765 | v3.0 |  |  | xvtstdcsp | VSX Vector Test Data Class SP |
| $111100 \ldots . .00000 \ldots . .11101$ 1011.. | XX2 | I | 766 | v3.0 |  |  | xvxexpdp | VSX Vector Extract Exponent DP |
| 111100 $\ldots . .01000 \ldots . .111011011 .$. | XX2 | I | 766 | v3.0 |  |  | xvxexpsp | VSX Vector Extract Exponent SP |
| 111100 $\ldots . .000001 \ldots . .111011011 .$. | XX2 | - | 767 | v3.0 |  |  | xvxsigdp | VSX Vector Extract Significand DP |
|  | XX2 | I | 767 | v3.0 |  |  | xvxsigsp | VSX Vector Extract Significand SP |
| 111100 $\ldots . . .10111 \ldots . .111011011 .$. | XX2 | I | 768 | v3.0 |  |  | xxbrd | VSX Vector Byte-Reverse Dword |
| 111100 $\ldots . .00111 \ldots . .111011011$. | XX2 | I | 768 | v3.0 |  |  | xxbrh | VSX Vector Byte-Reverse Hword |
| 111100 $\ldots . . .11111 \ldots . . .111011011 .$. | XX2 | I | 769 | v3.0 |  |  | xxbrq | VSX Vector Byte-Reverse Qword |
| 111100 $\ldots . . .01111 \ldots \ldots .111011011$. | XX2 | I | 769 | v3.0 |  |  | xxbrw | VSX Vector Byte-Reverse Word |
| 111100 $\ldots . .1 \ldots . . . . . .010100101 .$. | XX2 | I | 770 | v3.0 |  |  | xxextractuw | VSX Vector Extract Unsigned Word |
| 111100 $\ldots . .1 \ldots . . . . . .1010110101 .$. | XX2 | I | 770 | v3.0 |  |  | xxinsertw | VSX Vector Insert Word |
| 111100 $\ldots \ldots . . . . . . . . . .10000010 \ldots$ | XX3 | , | 771 | v2.06 |  |  | xxland | VSX Vector Logical AND |
| 111100 $\ldots . . .1 . . . . . . . . .10001010 \ldots$ | XX3 | - | 771 | v2.06 |  |  | xxlandc | VSX Vector Logical AND with Complement |
| 111100 $\ldots \ldots . \ldots \ldots . . . . . .10111010 \ldots$ | XX3 | I | 772 | v2.07 |  |  | xxleqv | VSX Vector Logical Equivalence |
| 111100 $\ldots \ldots . \ldots . . . . . . .10110010 \ldots$ | XX3 | I | 772 | v2.07 |  |  | xxInand | VSX Vector Logical NAND |
|  | XX3 | I | 773 | v2.06 |  |  | xxInor | VSX Vector Logical NOR |
| 111100 $\ldots \ldots . \ldots \ldots . . . . . . .10010010 \ldots$ | XX3 | I | 774 | v2.06 |  |  | xxlor | VSX Vector Logical OR |
| 111100 $\ldots \ldots . \ldots \ldots . . . . . . .10101010 \ldots$ | XX3 | 1 | 773 | v2.07 |  |  | xxlorc | VSX Vector Logical OR with Complement |
| 111100 $\ldots \ldots . . . . . . . . . . .100111010 \ldots$ | XX3 | I | 774 | v2.06 |  |  | xxlxor | VSX Vector Logical XOR |
|  | XX3 | I | 775 | v2.06 |  |  | xxmrghw | VSX Vector Merge Word High |
| $111100 \ldots . . . . . . . . . .00110010 \ldots$ | XX3 | I | 775 | v2.06 |  |  | xxmrglw | VSX Vector Merge Word Low |
| 111100 $\ldots . . . \ldots . . . . . . .00011010 \ldots$ | XX3 | I | 776 | v3.0 |  |  | xxperm | VSX Vector Permute |
| 11100 $\ldots . . . . . . . . . ~$ $0.01010 \ldots$ <br> 11100 $\ldots$  | XX3 | 1 | 777 | v2.06 |  |  | xxpermdi | VSX Vector Dword Permute Immediate |
| 111100 $\ldots . . . . . . . . . . . .00111010 \ldots$ | XX3 | - | 776 | v3.0 |  |  | xxpermr | VSX Vector Permute Right-indexed |
|  | XX4 | I | 777 | v2.06 |  |  | xxsel | VSX Vector Select |
|  | XX3 | I | 778 | v2.06 |  |  | xxsldwi | VSX Vector Shift Left Double by Word Immediate |
| 111100 $\ldots \ldots .00 \ldots \ldots 0.101101000$. | XX1 | 1 | 778 | v3.0 |  |  | xxspltib | VSX Vector Splat Immediate Byte |
| 111100 .... $111 . .1 . . .010100100$. | XX2 | I | 778 | v2.06 |  |  | xxspltw | VSX Vector Splat Word |

Figure 89. Power ISA Instruction Set Sorted by Mnemonic (Sheet 17 of 17)

1. Key to Instruction column (primary and extended opcode bits shaded in gray).

1 Instruction bit that corresponds to a reserved field, must have a value of 0 , otherwise invalid form.

- Instruction bit that corresponds to an operand bit, may have a value of either 0 or 1.
$0 \quad$ Instruction bit having a value 0.
$1 \quad$ Instruction bit having a value 1.

2. Key to Version column.

| P1 | Instruction introduced in the POWER Architecture. |
| :---: | :--- |
| P2 | Instruction introduced in the POWER2 Architecture. |
| PPC | Instruction introduced in the PowerPC Architecture prior to v2.00. |
| v2.00 | Instruction introduced in the PowerPC Architecture Version 2.00. |
| v2.01 | Instruction introduced in the PowerPC Architecture Version 2.01. |
| v2.02 | Instruction introduced in the PowerPC Architecture Version 2.02. |
| v2.03 | Instruction introduced in the Power ISA Architecture Version 2.03. |
| v2.04 | Instruction introduced in the Power ISA Architecture Version 2.04. |
| v2.05 | Instruction introduced in the Power ISA Architecture Version 2.05. |
| v2.06 | Instruction introduced in the Power ISA Architecture Version 2.06. |
| v2.07 | Instruction introduced in the Power ISA Architecture Version 2.07. |
| v3.0 | Instruction introduced in the Power ISA Architecture Version 3.00. |

3. Key to Privilege column.

P Denotes an instruction that is treated as privileged.
O Denotes an instruction that is treated as privileged or nonprivileged (or hypervisor, for mtspr), depending on the SPR or PMR number.
PI Denotes an instruction that is illegal in privileged state.
H Denotes an instruction that can be executed only in hypervisor state
U Denotes an instruction that can be executed only in ultravisor state
4. Key to Mode Dependency column.

Except as described below and in Section 1.11.3, "Effective Address Calculation", in Book I, all instructions are independent of whether the processor is in 32-bit or 64-bit mode.

[^18]
## Last Page - End of Document

Version 3.0


[^0]:    VRT $\leftarrow{ }^{96} 0 \|$ (VSCR)

[^1]:    Programming Note
    FX is defined not to be altered implicitly by mtfsfi and mtfsf because permitting these instructions to alter $F X$ implicitly can cause a paradox. An example is an $\boldsymbol{m t f s f i}$ or mtfsf instruction that supplies 0 for FX and 1 for $0 X$, and is executed when $0 X=0$. See also the Programming Notes with the definition of these two instructions.

    33 Floating-Point Enabled Exception Summary (FEX)
    This bit is the OR of all the floating-point exception bits masked by their respective enable bits. mcrfs, mtfsfi, mtfsf, mtfsbO, and $\boldsymbol{m t f s} \boldsymbol{b 1}$ cannot alter FEX explicitly.

[^2]:    1. Floating-point division is based on exponent subtraction and division of the significands.
    2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
[^3]:    1. Floating-point division is based on exponent subtraction and division of the significands.
    2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
[^4]:    1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
    2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( $\mathrm{G}, \mathrm{R}$, and X ) enter into the computation.
    3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
[^5]:    1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
    2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( $\mathrm{G}, \mathrm{R}$, and X ) enter into the computation.
    3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
[^6]:    1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
    2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
[^7]:    1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
    2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
[^8]:    1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
    2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( $\mathrm{G}, \mathrm{R}$, and X ) enter into the computation.
    3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
[^9]:    1. Floating-point multiplication is based on exponent addition and multiplication of the significands.
    2. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( $\mathrm{G}, \mathrm{R}$, and X ) enter into the computation.
    3. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
[^10]:    1. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( $G, R$, and $X$ ) enter into the computation.
    2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
[^11]:    Explanation:

    | src1 | The double-precision floating-point value in doubleword element $i$ of VSR[XA] (where $i \in\{0,1\})$. |
    | :--- | :--- |
    | src2 | The double-precision floating-point value in doubleword element $i$ of VSR[XB] (where $i \in\{0,1\})$. |
    | dQNaN | Default quiet NaN (0x7FF8_0000_0000_0000). |
    | NZF | Nonzero finite number. |
    | Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). |
    | $S(x, y)$ | Return the normalized sum of floating-point value $x$ and negated floating-point value $y$, having unbounded range and precision. |
    |  | Note: If $x=-y, v$ is considered to be an exact-zero-difference result (Rezd). |
    | $Q(x)$ | Return a QNaN with the payload of $x$. |

[^12]:    1. Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits ( $G, R$, and $X$ ) enter into the computation.
    2. Floating-point normalization is based on shifting the significand left until the most-significant bit is 1 and decrementing the exponent by the number of bits the significand was shifted.
[^13]:    Explanation:

    | src1 | The single-precision floating-point value in word element i of VSR[XA] (where i $\in\{0,1,2,3\}$ ). |
    | :---: | :---: |
    | src2 | The single-precision floating-point value in word element i of VSR[XB] (where $i \in\{0,1,2,3\}$ ). |
    | dQNaN | Default quiet NaN (0x7FC0_0000). |
    | NZF | Nonzero finite number. |
    | Rezd | Exact-zero-difference result (addition of two finite numbers having same magnitude but different signs). |
    | $\mathrm{S}(\mathrm{x}, \mathrm{y})$ | Return the normalized sum of floating-point value x and negated floating-point value y , having unbounded range and precision. Note: If $x=-y$, $v$ is considered to be an exact-zero-difference result (Rezd). |
    | $Q(x)$ | Return a QNaN with the payload of x . |
    | $v$ | The intermediate result having unbounded signficand precision and unbounded exponent range. |

    Table 138.Actions for xvsubsp

    ## VSX Vector Test for software Divide Double-Precision XX3-form

    

    Let $X A$ be the value $32 \times A X+A$.
    Let $X B$ be the value $32 \times B X+B$.
    $\mathrm{fe}_{\mathrm{e}} \mathrm{fl} \mathrm{ag}$ is initialized to 0 .
    $\mathrm{fg}_{\mathrm{-}}^{-} \mathrm{fl} \mathrm{ag}$ is initialized to 0 .
    For each vector element $i$ from 0 to 1 , do the following. Let srcl be the double-precision floating-point operand in doubleword element $i$ of VSR[ XA].

    Let $\mathrm{srCl}_{2}$ be the double-precision floating-point operand in doubleword element $i$ of VSR[ XB].

    Let $e_{-}$a be the unbiased exponent of src .
    Let $e_{-} b$ be the unbiased exponent of $\operatorname{src} 2$.
    fe_flag is set to 1 for any of the following conditions.

    - srcl is a NaN or an infinity.
    - srcl is a zero, a NaN, or an infinity.
    - e_b is less than or equal to 1022 .
    - $e_{-}^{-} b$ is greater than or equal to 1021 .
    - srcl is not a zero and the difference, $e_{-} a$ - $e_{-} b$, is greater than or equal to 1023.
    - srcl is not a zero and the difference, $e_{-} a \cdot e_{-} b$, is less than or equal to - 1021.
    - srcl is not a zero and e_a is less than or equal to - 970
    fg_flag is set to 1 for any of the following conditions.
    - srcl is an infinity.
    - $\operatorname{srcl}$ is a zero, an infinity, or a denormalized value.
    $C R$ field $B F$ is set to the value Ob1 ||fg_flag ||fe_flag || 0 bo.

    Special Registers Altered
    CR[BF]

    VSR Data Layout for xvtdivdp
    srcl = VSR[XA]

    | .$d$ word[0] | .$d$ word[1] |
    | :--- | :--- |

    $\operatorname{src2}=\operatorname{VSR}[X B]$
    

    ## Version 3.0

    ## VSX Vector Test for software Divide Single-Precision XX3-form

    ```
    xvtdivsp BF,XA,XB
    \begin{tabular}{|c|c|l|l|l|l|l|l|}
    \hline 60 & BF & II & & A & & B & \\
    \hline
    \end{tabular}
    \begin{tabular}{ll}
    \(X A\) & \(\leftarrow A X \| A\) \\
    \(X B\) & \(\leftarrow B X \| B\) \\
    eq_flag & \(\leftarrow 0\) b0 \\
    gt_flag & \(\leftarrow 0 b 0\)
    \end{tabular}
    do i=0 to 127 by 32
    src1 ```

[^14]:    loop: ldarx r1 $\leftarrow$ PTE_dwd_0 /* load dwd 0 of PTE */ if $\mathrm{V}=0$ abort, else/*to interact with locking*/ r157:60 $\leftarrow$ new SW value $/ *$ replace SW, in r1 */

[^15]:    STE $_{\text {ESID, }} \leftarrow$ new values ( $\mathrm{V}=1$ )
    ptesync /* order update before slbieg and before next Segment Table search */
    slbieg(old_B,old_ESID,old_TA,old_C)
    /*invalidate old translation*/
    eieio /* order slbieg before slbsync */
    slbsync /* order slbieg before ptesync */
    ptesync /* order slbieg, slbsync, and update before next data access

[^16]:    I Programming Note
    The PMCC field does not affect the behavior of the privileged Performance Monitor registers (SPRs 784-792, 795-798); accesses to these SPRs in problem state result in Privileged Instruction type Program interrupts.
    The PMCC field also does not affect the behavior of write operations to group B; write operations to SPRs in group B are treated as not supported regardless of privilege state. See the mtspr instruction description in Section 4.4.5 for additional information on accessing SPRs that are not supported.

[^17]:    CT If the instruction tests the Count Register, it tests the low-order 32 bits in 32-bit mode and all 64 bits in 64-bit mode.
    SR The setting of status registers (such as XER and CRO) is mode-dependent.
    32 The instruction can be executed only in 32-bit mode.
    64 The instruction can be executed only in 64-bit mode.

[^18]:    CT If the instruction tests the Count Register, it tests the low-order 32 bits in 32-bit mode and all 64 bits in 64-bit mode.
    SR The setting of status registers (such as XER and CRO) is mode-dependent.
    32 The instruction can be executed only in 32-bit mode.
    64 The instruction can be executed only in 64-bit mode.

