The math of Nxt forging

mthcl[*]

April 23, 2014. Version 0.4.2

Abstract

We discuss the forging algorithm of Nxt from the probabilistic point of view, and obtain explicit formulas and estimates for several important quantities, such as the probability that an account generates a block, the length of the longest sequence of consecutive blocks generated by one account, and the probability that one concurrent blockchain wins over another one.

1   Forging algorithm

In this article we concentrate on the 1-block-per-minute regime, which is not implemented yet. Assume that there are N forging accounts at a given (discrete) time, $B_1, \ldots, B_N$ are the corresponding balances, and we denote by

$$b_k = \frac{B_k}{B_1 + \cdots + B_N}, \quad k = 1, \ldots, N$$

the proportion of total forging power that the kth account has. Then, to determine which account will generate the next block, we take i.i.d. random variables with Uniform distribution on interval $(0, 1)$, and the account which maximizes b_k/U_k generates the block; i.e., the label k_0 of the generating account is determined by

$$k_0 = \operatorname*{arg\,min}_{j \in \{1, \ldots, N\}} \frac{U_j}{b_j}. \tag{1}$$

[*]NXT: 5978778981551971141; author's contact information: `e.monetki@gmail.com`, or send a PM at `bitcointalk.org` or `nxtforum.org`

We refer to the quantity U_k/b_k as the *weight* of the kth account, and to U_{k_0}/b_{k_0} as the weight of the block. This procedure is called the *main algorithm* (because it is actually implemented in Nxt at this time), or the *U-algorithm* (because the weights have Uniform distribution).

As a general rule, it is assumed that the probability that an account generates a block is proportional to the effective balance, but, in fact, this is only approximately true (as we shall see in Section 2). For comparison, we consider also the following rule of choosing the generating account: instead of (1), we use

$$k_0 = \arg\min_{j \in \{1,\dots,N\}} \frac{|\ln U_j|}{b_j}. \tag{2}$$

The corresponding algorithm is referred to as *Exp-algorithm* (since the weights now have Exponential probability distribution).

Important note: for all the calculations in this article, we assume that all accounts are forging and all balances are effective (so that $B_1 + \cdots + B_N$ equals the total amount of NXT in existence). In the real situation when only the proportion α of all money is forging, one can adjust the formulas in the following way. Making a reasonable assumption that *all* money of the bad guy is forging and his (relative) stake is b', all the calculations in this article are valid with $b = \alpha^{-1} b'$.

2 Probability of block generation

Observe that (see e.g. Example 2a of Section 10.2.1 of [5]) the random variable $|\ln U_j|/b_j$ has Exponential distribution with rate b_j. By (2), the weight of the generated block is also an Exponential random variable with rate $b_1 + \cdots + b_N = 1$ (cf. (5.6) of [6]), and the probability that the kth account generates the block is exactly b_k (this follows e.g. from (5.5) of [6]).

However, for U-algorithm the calculation in the general case is not so easy. We concentrate on the following situation, which seems to be critical for accessing the security of the system: N is large, the accounts $2, \dots, N$ belong to "poor honest guys" (so $b_2, \dots, b_N$ are small), and the account 1 belongs to a "bad guy", who is not necessarily poor (i.e., $b := b_1$ need not be very small).

We first calculate the probability distribution of the biggest weight among the good guys: for $x \gg \max_{k \geq 2} b_k$ let us write

$$
\begin{aligned}
\mathbb{P}\left[\max_{k \geq 2} \frac{b_k}{U_k} < x\right] &= \prod_{k \geq 2} \mathbb{P}\left[U_k > \frac{b_k}{x}\right] \\
&= \prod_{k \geq 2}\left(1 - \frac{b_k}{x}\right) \\
&= \exp \sum_{k \geq 2} \ln\left(1 - \frac{b_k}{x}\right) \\
&\approx e^{-\frac{1-b}{x}},
\end{aligned} \tag{3}
$$

since $\ln(1-y) \sim -y$ as $y \to 0$ and $b_2 + \cdots + b_N = 1 - b$. We calculate now the probability $f(b)$ that the bad guy generates the block, in the following way. Let Y be a random variable with distribution (3) and independent of U_1, and we write (conditioning on U_1)

$$
\begin{aligned}
f(b) &:= \mathbb{P}\left[\frac{b}{U_1} > Y\right] \\
&= \int_0^1 \mathbb{P}\left[Y < \frac{b}{z}\right] dz \\
&= \int_0^1 e^{-\frac{1-b}{b}z} dz \\
&= \frac{b}{1-b}\left(1 - e^{-\frac{1-b}{b}}\right).
\end{aligned} \tag{4}
$$

It is elementary to show that $f(b) > b$ for all $b \in (0,1)$ (see also Figure 1), and (using the Taylor expansion) $f(b) = b + b^2 + O(b^3)$ as $b \to 0$.

Let us remark also that, since $f(b) \sim b$ as $b \to 0$ and using a calculation similar to (3), if *all* relative balances are small, the situation very much resembles that under Exp-algorithm (see also (9) below).

2.1 Splitting of accounts

Here we examine the situation when an owner of an account splits it into two (or even several) parts, and show that, in general, this strategy is not favorable to the owner.

Figure 1: The plot of $f(b)$ (black curve)

First of all, as discussed in the beginning of Section 2, for the Exp-algorithm, the probability that one of the new (i.e., obtained after the splitting) accounts will generate the next block does not change at all. Indeed, this probability is exactly the proportion of the total balance owned by the account, and any splitting does not change this proportion (i.e., all the new accounts forge *exactly* as the old one).

Now, let us consider the case of U-algorithm. We shall prove that splitting is *always* unfavorable for a Nxt holder. Namely, we can prove the following result[1]:

Theorem 2.1. *Assume that a person or entity controls a certain number of Nxt accounts, and let p be the probability of generating the next block (i.e., the account that forges the block belongs to this person or entity). Suppose now that one of these accounts is split into two parts (while the balances of all other accounts stay intact), and let p' be the probability of block generation in this new situation. Then $p' < p$.*

By induction, one easily obtains the following

Corollary 2.2. *Under the U-algorithm, in order to maximize the probability of generating the next block, all NXT that one controls should be concentrated in only one account.*

Proof of Theorem 2.1. Let $b_1, \ldots, b_\ell$ be the relative balances of accounts controlled by that person or entity, and let $b_{\ell+1}, \ldots, b_n$ be the balances of the other active accounts. Assume without restriction of generality that the first account is split into two parts with (positive) relative balances b_1', b_1'' (so that $b_1' + b_1'' = b_1$).

Let $U_1, \ldots, U_n, U_1', U_1''$ be i.i.d. Uniform$[0,1]$ random variables. Let

$$Y = \min_{j=1,\ldots,\ell} \frac{U_j}{b_j},$$

$$Y' = \min\left(\frac{U_1'}{b_1'}, \frac{U_1''}{b_1''}, \min_{j=2,\ldots,\ell} \frac{U_j}{b_j}\right),$$

$$Z = \min_{j=\ell+1,\ldots,n} \frac{U_j}{b_j}.$$

---

[1]The author is happy that he is able to add at least one theorem to this text. Without theorems, he had a strong feeling of doing something *unusual*.

Let us denote $x^+ := \max(0, x)$ for $x \in \mathbb{R}$. Analogously e.g. to (3), we have for $t > 0$

$$\mathbb{P}[Y > t] = \prod_{j=1,\dots,\ell} (1 - b_j t)^+,$$

$$\mathbb{P}[Y' > t] = (1 - b_1' t)^+ (1 - b_1'' t)^+ \prod_{j=2,\dots,\ell} (1 - b_j t)^+,$$

and a similar formula holds for Z; we, however, do not need the explicit form of the distribution function of Z, so we just denote this function by ζ.

Observe that for $0 < t < \min\left(\frac{1}{b_1'}, \frac{1}{b_1''}\right)$ it holds that

$$(1 - b_1' t)(1 - b_1'' t) = 1 - b_1 t + b_1' b_1'' t^2$$
$$> 1 - b_1 t,$$

so

$$(1 - b_1' t)^+ (1 - b_1'' t)^+ \geq (1 - b_1 t)^+$$

for all $t \geq 0$ (if the left-hand side is equal to 0, then so is the right-hand side), and, moreover, the inequality is strict in the interval $\left(0, \min\left(\frac{1}{b_1'}, \frac{1}{b_1''}\right)\right)$.

Then, conditioning on Z, we obtain

$$1 - p = \mathbb{P}[Y > Z]$$
$$= \int_0^\infty \prod_{j=1,\dots,\ell} (1 - b_j t)^+ \, d\zeta(t)$$
$$= \int_0^\infty (1 - b_1 t)^+ \prod_{j=2,\dots,\ell} (1 - b_j t)^+ \, d\zeta(t)$$
$$< \int_0^\infty (1 - b_1' t)^+ (1 - b_1'' t)^+ \prod_{j=2,\dots,\ell} (1 - b_j t)^+ \, d\zeta(t)$$
$$= \mathbb{P}[Y' > Z]$$
$$= 1 - p',$$

and this concludes the proof of the theorem. $\qquad\square$

One should observe, however, that the disadvantage of splitting under the U-algorithm is not very significant. For example, if one person controls 5% of total active balance and has only one account, then, according to (4), the forging probability is approximately 0.0526. For *any* splitting, this probability cannot fall below 0.05 (this value corresponds to the the extreme situation when all this money is distributed between many small accounts).

Conclusions:

1. Under Exp-algorithm, the probability that an account with relative active balance b generates the next block is exactly b; if all relative balances are small, then the U-algorithm essentially works the same way as the Exp-algorithm.

2. For the U-algorithm, if an account has proportion b of the total active balance and the forging powers of other accounts are relatively small, then the probability that it generates the next block is given by $f(b)$ of (4).

3. With small b, $f(b) \approx b + b^2$, i.e., the block generating probability is roughly proportional to the effective balance with a quadratic correction.

4. It is also straightforward to obtain that the probability that a good guy k generates a block is $b_k(1 - f(b))$, up to terms of smaller order.

5. In general, splitting has no effect on the (total) probability of block generation under Exp-algorithm, and this probability always decreases under U-algorithm. However, the difference is usually not very significant (even if the account is split to many small parts).

6. Thus, neither algorithm encourages splitting (anyhow, there is some cost in maintaining many forging accounts, so, in principle, there is no reason to increase too much the number of them in the case of Exp-algorithm as well). The reader should be warned, however, that all the conclusions in this article are valid for *mathematical models*, and the real world can introduce some corrections.

7. In particular, it should be observed that, if the attacker could harm the network by splitting his account into many small ones, then a very small gain that he achieves by not splitting would not prevent him from attacking the network. If this attacker's strategy presents any real danger, we may consider introducing a *lower limit* for forging (e.g., only accounts with more than, say, 100 NXT are allowed to forge).

3 Longest run

We consider a "static" situation here: there are no transactions (so that the effective balances are equal to full balances and do not change over time). The goal is to be able to find out, how many blocks in a row can be typically generated by a given account over a long period n of time.

So, assume that the probability that an account generates the next block is p (see in Section 2 an explanation about how p can be calculated). It is enough to consider the following question: let R_n be the maximal number of consecutive 1's in the sequence of n Bernoulli trials with success probability p; what can be said about the properties of the random variable R_n?

The probability distribution of R_n has no tractable closed form, but is nevertheless quite well studied, see e.g. [7] (this article is freely available in the internet). The following results are taken from [4]: we have

$$\mathbb{E}R_n = \log_{1/p} qn + \frac{\gamma}{\ln 1/p} - \frac{1}{2} + r_1(n) + \varepsilon_1(n), \tag{5}$$

$$\mathrm{Var}R_n = \frac{\pi^2}{6\ln^2 1/p} + \frac{1}{12} + r_2(n) + \varepsilon_2(n), \tag{6}$$

where $q = 1 - p$, $\gamma \approx 0.577\ldots$ is the Euler-Mascheroni constant, $\varepsilon_{1,2}(n) \to 0$ as $n \to \infty$, and $r_{1,2}(n)$ are uniformly bounded in n and very small (so, in practice, $r_{1,2}$ and $\varepsilon_{1,2}$ can be neglected).

In the same work, one can also find results on the distribution itself. Let W_p be a random variable with Gumbel-type distribution: for $y \in \mathbb{R}$

$$\mathbb{P}[W_p \le y] = \exp(-p^{y+1}).$$

Then, for $x = 0, 1, 2, \ldots$ it holds that

$$\mathbb{P}[R_n = x] \approx \mathbb{P}[x - \log_{1/p} qn < W_p \le x + 1 - \log_{1/p} qn], \tag{7}$$

with the error decreasing to 0 as $n \to \infty$. So, in particular, one can obtain that

$$\mathbb{P}[R_n \ge x] \approx 1 - \exp(-p^{x+1} qn)$$
$$\approx p^{x+1} qn \tag{8}$$

if $p^{x+1}qn$ is small (the last approximation follows from the Taylor expansion for the exponent).

For example, consider the situation when one account has 10% of all forging power, and the others are relatively small. Then, according to (4), the probability that this account generates a block is $p \approx 0.111125$. Take $n = 1000000$, then, according to (5)–(7), we have

$$\mathbb{E}R_n \approx 6.00273,$$
$$\mathrm{Var}R_n \approx 0.424,$$
$$\mathbb{P}[R_n \geq 7] \approx 0.009 \,.$$

Conclusions:

1. The distribution of the longest run of blocks generated by one particular account (or group of accounts) is easily accessible, even though there is no exact closed form. Its expectation and variance are given by (5)–(6), and the one-sided estimates are available using (8).

4   Weight of the blockchain and concurrent blockchains

First, let us look at the distribution of the weight of a block. In the case of Exp-algorithm, everything is simple: as observed in Section 2, it has the Exponential distribution with rate 1. This readily implies that the expectation of the sum of weights of n blocks equals n.

As for the U-algorithm, we begin by considering the situation when all relative balances are small. Analogously to (3), being W the weight of the block, for $x \ll (\max_k b_k)^{-1}$ we calculate

$$\begin{aligned}
\mathbb{P}\Big[\frac{1}{W} > x\Big] &= \mathbb{P}\Big[\max_k \frac{b_k}{U_k} < \frac{1}{x}\Big] \\
&= \prod_k \mathbb{P}\Big[U_k > xb_k\Big] \\
&= \prod_k (1 - xb_k) \\
&= \exp \sum_k \ln(1 - xb_k) \\
&\approx e^{-x},
\end{aligned} \tag{9}$$

so also in this case the distribution of the weight is approximately Exponential with rate 1.

We consider now the situation when all balances except the first one are small, and $b := b_1$ need not be small. For the case of U-algorithm, similarly to (9) we obtain for $x \in (0, 1/b)$

$$
\mathbb{P}\left[\frac{1}{W} > x\right] = \prod_k (1 - xb_k)
$$
$$
= (1 - bx) \exp \sum_{k \geq 2} \ln(1 - xb_k)
$$
$$
\approx (1 - bx)e^{-(1-b)x}, \tag{10}
$$

so

$$
\mathbb{E}\frac{1}{W} \approx \int_0^{1/b} (1 - bx)e^{-(1-b)x} \, dx
$$
$$
= \frac{be^{-\frac{1-b}{b}} + 1 - 2b}{(1 - b)^2}. \tag{11}
$$

One can observe (see Figure 2) that the right-hand side of (11) is strictly between $1/2$ and 1 for $b \in (0, 1)$.

Let us consider now the following attack scenario: account 1 (the "bad guy", with balance b) temporarily disconnects from the network and forges its own blockchain; he then reconnects hoping that his blockchain would be "better" (i.e., has smaller sum of weights). Then, while the account 1 is disconnected, the "good" part of the network produces blocks with weights having Exponential distribution with rate $1 - b$, and thus each weight has expected value $\frac{1}{1-b}$.

Let $X_1, X_2, X_3, \ldots$ be the weights of the blocks produced by the "good part" of the network (after the bad guy disconnects), and we denote by $Y_1, Y_2, Y_3, \ldots$ the weights of the blocks produced by the bad guy. We are interested in controlling the probability of the following event (which means that the blockchain produced by the bad guy has smaller sum of weights and therefore wins)

$$
H_m = \{X_1 + \cdots + X_m - Y_1 - \cdots - Y_m \geq 0\}
$$

for "reasonably large" m (e.g., $m = 10$ or so). If the probability of H_m is small, this means that the bad guy just does not have enough power to

10

Figure 2: Expectation of the weight (as a function of b)

attack the network; on the other hand, if this probability is not small, then the system should be able to fence off the attack by other means, which we shall not discuss in this note.

We obtain an upper bound on the probability of the event H_m using the so-called Chernoff theorem (see e.g. Proposition 5.2 of Chapter 8 of [5]). More specifically, using Chebyshev's inequality and the fact that the random variables are i.i.d., we write for any fixed $t > 0$

$$
\begin{aligned}
\mathbb{P}[H_m] &= \mathbb{P}[X_1 + \cdots + X_m - Y_1 - \cdots - Y_m \geq 0] \\
&= \mathbb{P}\big[\exp\big(t(X_1 + \cdots + X_m - Y_1 - \cdots - Y_m)\big) \geq 1\big] \\
&\leq \mathbb{E}\exp\big(t(X_1 + \cdots + X_m - Y_1 - \cdots - Y_m)\big) \\
&= \big(\mathbb{E}e^{t(X_1-Y_1)}\big)^m.
\end{aligned}
$$

Since the above is valid for all $t > 0$, we have

$$
\mathbb{P}[H_m] \leq \delta^m, \quad \text{where} \quad \delta = \inf_{t>0} \mathbb{E}e^{t(X_1-Y_1)}. \tag{12}
$$

It is important to observe that this bound is nontrivial (i.e., $\delta < 1$) only in the case $\mathbb{E}X_1 < \mathbb{E}Y_1$.

For U-algorithm, X_1 is Exponentially distributed with rate $1 - b$, and Y_1 has Uniform$(0, b^{-1})$ distribution. So, the condition $\mathbb{E}X_1 < \mathbb{E}Y_1$ is equivalent to $(1 - b)^{-1} < (2b)^{-1}$, that is, $b < 1/3$. Then, for $b < 1/3$, the parameter δ from (12) is determined by

$$
\delta = \delta(b) = b(1-b) \inf_{0<t<1-b} \frac{1 - e^{-t/b}}{t(1 - b - t)} \tag{13}
$$

(see the plot of $\delta(b)$ on Figure 3), so we have

$$
\mathbb{P}[H_m] \leq \delta(b)^m. \tag{14}
$$

For example, for $b = 0.1$ we have $\delta(b) \approx 0.439$. We have, however, $\delta(b) \approx 0.991$ for $b = 0.3$, which means that one has to take very large m in order to make the right-hand side of (14) small in this case.

For the Exp-algorithm, the bad guy would produce blocks with weights having Exponential distribution with rate b, so each weight has expected value $\frac{1}{b}$. Similarly to the above, one obtains that the condition $\mathbb{E}X_1 < \mathbb{E}Y_1$ is equivalent to $b < 1/2$, and

$$
\mathbb{P}[H_m] \leq \big(4b(1 - b)\big)^m \tag{15}
$$

(that is, δ can be explicitly calculated in this case and equals $4b(1 - b)$; observe that $4b(1 - b) < 1$ for $b < 1/2$).

Figure 3: The plot of $\delta(b)$

13

Conclusions:

1. We analyze an attack strategy when one account (or a group of accounts) temporarily disconnects from the main network and tries to forge a "better" blockchain than the one forged by other accounts, in the situation when one bad rich guy has proportion b of total amount of NXT, and the stakes of the others are relatively small.

2. The probability that the bad guy forges a better chain of length m can be controlled using (14) (for the U-algorithm) or (15) (for the Exp-algorithm).

3. It should be observed that this probability does not tend to 0 (as $m \to \infty$) if the bad guy has at least 1/3 of all *active* balances in the network in the case of U-algorithm (correspondingly, at least 1/2 in the case of Exp-algorithm). There should exist some specific methods for protecting the network against such an attack in the case when there is risk that (active) relative balance of the bad guy could become larger than the above threshold.

4. For the current realization of the U-algorithm, the author expects that this analysis can be performed in a quite similar way (because the weight is then proportional to the time to the next block, and the longest blockchain wins), with an additional difficulty due to the oscillating `BaseTarget`.

5. It may be a good idea to limit the forging power of accounts by some fixed threshold, e.g., if an account has more than, say, 1M NXT, then it forges as if it had exactly 1M NXT. Of course, a rich guy can split his fortune between smaller accounts, but then all those accounts would forge roughly as one big account (without threshold) under Exp-algorithm. So, one can use the computationally easier U-algorithm without having its drawbacks (the 1/3 vs. 1/2 issue) discussed in this section.

5 More on account splitting

In this section we analyze the following attack strategy: if the bad guy wins the forging lottery at the current step but owns *several* accounts with weights

less than all the accounts of good guys (i.e., these accounts of the bad guy are first in the queue), then he chooses which of his accounts will forge. Effectively, that means that he has several independent tries for choosing the hash of the block, and so he may be able to manipulate the lottery in his favor.

Of course, following this strategy requires that the balance of the winning accounts should be small (because of the ban for non-forging), but, as we shall see below, splitting into small parts is exactly the right strategy for maximizing the number of the best accounts in the queue.

First of all, let us estimate the probability that in the sequence of accounts ordered by the weights, the first k_0 ones belong to the same person or entity, who controls the proportion b of all active balance. We will do the calculations for the case of Exp-algorithm, since the attacker would have to split his money between many (or at least several) accounts anyway, and, as observed in Section 2, in this situation both algorithms work essentially in the same way.

One obvious difficulty is that we do not know, how exactly the money of the attacker are distributed between his accounts. It is reasonable, however, to assume that the balances of the other accounts (those not belonging to the attacker) are relatively small. Let us show the following remarkable fact:

Proposition 5.1. *The best strategy for the attacker to maximize the probability that the best k_0 accounts belong to him (under the Exp-algorithm), is to split his money uniformly between many accounts.*

Proof. We assume that r is the *minimal* relative balance per account that is possible, and let us assume that the attackers money are held in accounts with relative balances $n_1 r, \ldots, n_\ell r$, where $\ell \geq k_0$. Denote also $m = n_1 + \cdots + n_\ell$, so that $b = mr$. Now, we make use of the elementary properties of the Exponential probability distribution discussed in the beginning of Section 2. Consider i.i.d. Exponential(r) random variables $Y_1, \ldots, Y_m$, and let $Y_{(1)}, \ldots, Y_{(k_0)}$ be the first k_0 order statistics of this sample. Then, abbreviating $s_j = n_1 + \cdots + n_j$, $s_0 := 0$, we have that

$$Z_j = \min_{i = s_{j-1}, \ldots, s_j} Y_i, \qquad j = 1, \ldots, \ell$$

are independent Exponential random variables with rates $n_1 r, \ldots, n_\ell r$. So, the orders statistics of Z-variables form a subset of the order statistics of Y-variables; since the weights of the attacker's accounts are the first k_0 order statistics of $Z_1, \ldots, Z_\ell$, the claim follows. $\square$

The above proposition leads to a surprisingly simple (approximate) upper bound for the probability that the best k_0 accounts belong to the attacker. Assume that *all* the accounts in the network have the minimum relative balance r; then each account in the (ordered with respect to the weights) sequence has probability b to belong to the attacker. Since k_0 should be typically small compared to the total number of accounts, we may assume that the ownerships of the first k_0 accounts are roughly independent, and this means that the probability that all the best k_0 accounts belong to the attacker should not exceed b^{k_0}.

Now, let us estimate the conditional probability $p^*(b)$ that the attacker wins the forging lottery on the next step, given he was chosen to forge at the current step. The above analysis suggests that the number of the best accounts in the queue that belong to the attacker can be approximated by a Geometric distribution with parameter b. Now, given that the attacker owns k best accounts in the queue, the probability that he wins the next forging lottery is $1 - (1 - b)^k$ (since there are k independent trials at his disposal, and the probability that all will fail is $(1 - b)^k$).

Using the memoryless property of the Geometric distribution (i.e., as one can easily verify, $\mathbb{P}[X = k] = \mathbb{P}[X = k + n \mid X \geq n]$ if X has the Geometric law) we have that, given that the winning account belongs to the attacker, he also owns the next $k - 1$ ones with probability $(1 - b)b^{k-1}$. So,

$$
\begin{aligned}
p^*(b) &= \sum_{k=1}^{\infty} (1 - b)b^{k-1}(1 - (1 - b)^k) \\
&= (1 - b)\sum_{j=0}^{\infty} b^j + (1 - b)^2 \sum_{j=0}^{\infty}(b(1 - b))^j \\
&= 1 - \frac{(1 - b)^2}{1 - b(1 - b)},
\end{aligned}
\tag{16}
$$

see the plot of the above function on Figure 4. The quantity $p^*(b)$ is almost b for small stakes (e.g., 0.1099 for $b = 0.1$), but dramatically increases for large b. For $b = 0.9$, for instance, this probability becomes 0.989, i.e., the attacker will be able to forge typically around 90 blocks in a row, instead of just 10.

Observe also, that this calculation applies to the following strategy of the attacker: take the *first* of his accounts that assures that he forges on the next step (so that the attacker minimizes the number of his accounts that

16

Figure 4: The plot of $p^*(b) = 1 - \frac{(1-b)^2}{1-b(1-b)}$.

will get banned). For this strategy, the attacker looks only one step to the future. One can consider also a more advanced strategy: since the future is (for now) deterministic, the attacker can try to calculate deeper into the future. We shall prove now that, under current implementation of the forging algorithm, *an attacker who owns more than 50% of all NXT can eventually forge all the blocks* (i.e., at some moment he starts forging and never stops). To prove this, we construct a *Galton-Watson* branching process (cf. e.g. [2] for the general theory) in the following way. Assume that the block $\ell_0 + 1$ is about to be forged. Let $Z_0 := 1$, and let Z_1 be the number of attacker's accounts that are first in the queue (i.e., win over any account not belonging to the attacker). Now, the attacker can choose which of these Z_1 accounts will forge. Let $Z_2^{(j)}$, $j = 1, \ldots, Z_1$ be the number of attacker's accounts that are first in the queue for the block $\ell_0 + 2$, provided he has chosen the jth account to forge at the previous step. Let $Z_2 = Z_2^{(1)} + \cdots + Z_2^{(Z_1)}$. Then, we define $Z_3, Z_4, Z_5, \ldots$ in an analogous way; it is then elementary to see that $(Z_n, n \geq 0)$ is a Galton-Watson branching process with the offspring law given by $p_k = (1 - b)b^k$, $k \geq 0$. The mean number of offspring

$$\mu = \sum_{k=1}^{\infty} k(1 - b)b^k = \frac{b}{1 - b}$$

is strictly greater than 1 when $b > \frac{1}{2}$. Since a supercritical branching process survives with positive probability (in fact, one can calculate that in this case the probability of survival equals $\frac{1-b}{b}$), the attacker can choose an infinite branch in the genealogical tree of the branching process, and follow it.

The attacker can also use the same strategy with $b \leq \frac{1}{2}$; this, of course, will not permit him to forge all the blocks, but there is still a possibility to increase the number of generated blocks. Let us do the calculations. The probability generating function of the number of offspring (corresponding to the Geometric distribution) is

$$g(s) = \sum_{j=0}^{\infty} s^j (1 - b)b^j = \frac{1 - b}{1 - sb},$$

so $a_n(b) := \mathbb{P}[Z_n = 0]$ satisfies the recursion

$$a_1(b) = 1 - b, \qquad a_{n+1}(b) = \frac{1 - b}{1 - ba_n(b)},$$

Figure 5: The plot of $\frac{h(b)}{1+h(b)}$ (it is not very evident from this picture, but $\frac{h(b)}{1+h(b)} \to 1$ as $b \to \frac{1}{2}$).

and the mean lifetime $h(b)$ of the branching process is

$$h(b) = \sum_{n=1}^{\infty}(1 - a_n(b)).$$

Unfortunately, usually there is no closed form expression for the expected lifetime of a subcritical branching process, but, as a general fact, it holds that $h(b) \sim b$ as $b \to 0$ and $h(b) \to \infty$ as $b \to \frac{1}{2}$. Since each streak of attacker's blocks has the expected length $b^{-1}h(b)$ and the expected length of each streak of good guy's blocks is b^{-1}, the attacker is able to forge the proportion $\frac{h(b)}{1+h(b)}$ of all blocks, see Figure 5 (the author thanks `Mathematica` for doing the computations).

Conclusions:

1. The probability that the attacker controls the best k_0 accounts can be bounded from above by b^{k_0}, where b is the attacker's stake.

19

2. Under current implementation of the forging algorithm, an attacker who owns more than 50% of all NXT can eventually forge all the blocks (i.e., at some moment he starts forging and never stops).

3. For additional network protection, we can recommend to have also a *lower limit* for forging, i.e., an account that has less than (say) 1000 NXT does not forge at all.

4. In fact, it may be a good idea to change this lower limit dynamically: In normal situation (there are not many non-forging events) it can be relatively low. However, if someone is starting playing games (and so there are many non-forging events), then the lower limit increases in order to protect the network. Incidentally, this increase of the lower limit will greatly decrease the attacking strength of the bad guy, since most of his accounts suddenly are unable to forge at all.

5. The community should be warned that if someone is advertising a forging pool but makes the forgers link to many different accounts, then it is a *very* suspicious behavior.

6 How to produce a random number?

All the previous discussion in this paper was restricted to the mathematical model of Section 1, defined in terms of i.i.d. Uniform random variables $U_1, \ldots, U_n$. This, however, is only an approximation of reality (similarly to all other mathematical models in this world); in fact, the U-variables are *pseudorandom*, i.e., they are deterministically computed from the information contained in the previous block together with the account hashes. That is, in the static situation (no transactions and the nodes do not connect/disconnect to/from the network) one can precisely determine who will generate the subsequent blocks. Some people have expressed concern whether this situation is potentially dangerous, i.e., an attacker could somehow exploit this in order to forge "too many" (= more than the probability theory permits) blocks in a row. To fence off this kind of treat, we may want to introduce some "true randomness" to the system; i.e., we want the network to produce a "truly unpredictable" random number U (say, with Uniform distribution on $[0, 1]$) in a decentralized way, with no central authority. Of course, one is not obliged to use this random number for forging (see the first remark in

"conclusions" below), but nevertheless the ability to produce random numbers may be useful for other applications, e.g., lotteries on top of the Nxt blockchain.

In principle, this is not an easy task, since the nodes controlled by the attacker can cheat by producing some carefully chosen nonrandom numbers, and it is not clear, how does the network recognize who cheats, and who does not. To deal with this problem, we first consider the following algorithm[2]:

- each account obtains a random number (say, in the interval $[0, 1]$) using some *local* randomizing device (e.g., `rand()` or whatever), and publishes the *hash* of this number;

- we calculate the weights of all accounts (using U-algorithm or Exp-algorithm);

- then, first k_0 accounts (with respect to the weights, in the increasing order) publish the numbers themselves;

- if at least one the published number does not correspond to its hash or at least one chosen account does not publish its number at all, the corresponding account is penalized (i.e., not allowed to forge during some time), and the whole procedure must be repeated (immediately, or somewhat later);

- we then "mix" these numbers (e.g., by summing them modulo 1^3), to obtain the random number we are looking for.

The parameter k_0 is supposed to be large enough so that the attacker would never control *exactly all* of k_0 best (with respect to the weights) accounts. Below we will discuss the question how k_0 should be chosen, depending on the maximal amount of active balance that the attacker can obtain. At this point it is important to observe that even one "honest" account in k_0 is enough; indeed, for the above mixing method, this follows from the fact that if U is a Uniform$[0, 1]$ random variable and X is *any* random variable independent of U, then $(X + U \mod 1)$ is also a Uniform$[0, 1]$ random variable.

Note that this two-step procedure (first publish the hash, and only then the number itself) is necessary. If we do not obscure the numbers, then the

---

[2]a similar procedure was proposed in [3].

[3]By definition, $(x \mod 1)$ is the fractional part of x, i.e., set the integer part to 0 and keep the digits after the decimal point.

attacker can see the $k_0 - 1$ numbers that are already published, and then publish something nonrandom that suits him. If we obscure them first, then the attacker cannot manipulate the procedure.

Let us explain also why the procedure should be restarted when at least one account *attempts* to cheat. It seems to be more "economical" to just pick the next account in the queue if some previous account excludes itself for whatever reason. However, this opens the door for the following attacking strategy. Assume, for example, that first $k_0 - 1$ accounts have already revealed their numbers, and the k_0th and $(k_0 + 1)$th accounts belong to the attacker. Then, he can actually *choose*, which of the two numbers will be published; this, obviously, creates a bias in his favor. In fact, we will see below that the best strategy for the attacker is to have many small accounts, so, invalidating one round of this procedure would not cost much to him. However, still each attempt costs one banned account, and, more importantly, if many rounds of the procedure are invalidated, it is likely that the identity of the attacker could be revealed (one can analyze the blockchain to investigate the origin of the money in offending accounts).

As discussed in Section 5, the probability that all the best k_0 accounts belong to the attacker can be bounded from above by b^{k_0}. For example, if $b = 0.9$ (i.e., the attacker has 90% of all NXT) and $k_0 = 150$, then $b^{k_0} \approx 0.00000013689$.

There are, however, questions about the practical implementation of this algorithm: the difficulty about obtaining consensus on who are the top accounts in the lottery, if such a procedure can slow down the network, etc. Therefore, we propose

Another randomization algorithm. It is described in the following way:

(a) each forging account must maintain a hashchain of some (large) length (actually, the nodes must maintain them for their accounts), and publish the last hash;

(b) at blocks $N\ell, \ell = 1, 2, 3, \ldots$, the list of K richest accounts with respect to the effective balance is formed, and this list becomes *valid* for blocks $N\ell + L, \ldots, N(\ell + 1) + L - 1$

(c) special "randomizing" blocks are forged just in between of every two normal blocks (this can be adjusted, maybe does not need to be so frequent);

(d) randomizing blocks are forged by the accounts from the valid list, e.g., in the cyclic order; they contain the hash from the hashchain preceding to the one already published (so the forger cannot cheat);

(e) to determine the forger of the next block (number n), the random number we use is `sha256(what was there before, sum of hashes in $L'$ last randomizing blocks published before the block $n-L$).`

One may take e.g. $N = 1440, K = 100, L = 10, L' = 50$, but, of course, these constants may be adjusted. In particular, the value of K must be neither too small, nor too large. If it is small, the danger is that the bad guy controls all K top accounts. On the other hand, if it is too large, then some accounts among the top K would be (relatively) small (for a rich guy), so he can start thinking about playing non-forging games.

With this approach, by the way, we can predict the next L forgers (but not more!), which was also a desirable feature of TF, a far as the author remembers.

The point of having the richest accounts do the randomization job is the following: it is probably impossible for the attacker to control e.g. top 100 accounts, that is just too expensive. And another point: since the accounts must be big, cheating by not publishing the random number now becomes very expensive as well (an account that does not forge the randomizing block when it must to, is banned and so unable to forge normal blocks for some period).

The contents of this section is based on private discussions with `ChuckOne` and `mczarnek` at `nxtforum.org`; see also [1].

Conclusions:

1. The current forging algorithm is only pseudorandom (deterministic but unpredictable), and there is concern whether this situation could be potentially dangerous. The author does think that this danger is not very serious, since the real world will not hesitate in introducing some "real randomness" to the system (because nodes go online and offline, money are transferred, etc.).

2. Nevertheless, it is possible to propose an extra randomization algorithm, i.e., the network can achieve a consensus on a Uniform[0,1] random number independent of the previously published data.

3. In this section we discuss two variants of such an algorithm.

Acknowledgments

References

[1] JFTR: randomness without cheating.
   `https://nxtforum.org/transparent-forging/jftr-randomness-without-cheating`

[2] K.B. ATHREYA, P.E. NEY (1972) *Branching Processes.* Springer-Verlag Berlin–Heidelberg–New York.

[3] MATTHEW CZARNEK. Transparent forging algorithm.
   `nxtforum.org/transparent-forging-*/transparent-forging-algorithm-245/`

[4] L. GORDON, M.F. SCHILLING, M.S. WATERMAN (1986) An extreme value theory for long head runs. *Probab. Theory Relat. Fields* **72**, 279–287.

[5] SHELDON M. ROSS (2009) *A First Course in Probability.* 8th ed.

[6] SHELDON M. ROSS (2012) *Introduction to Probability Models.* 10th ed.

[7] MARK SCHILLING (1990) The Longest Run of Heads. *The College Math J.,* **21** (3), 196–206.