LEDGER
ledgerjournal.org

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.46

A Probabilistic Analysis of the Nxt Forging Algorithm: Open Review

Author: Serguei Popov*†

Reviewers: Reviewer A, Reviewer B, Reviewer C

Abstract. The final version of the paper "Gaming Self-Contained Provably Fair Smart Contract Casinos" can be found in Ledger Vol. 1 (2016) 69-83, DOI 10.5915/LEDGER.2016.46. There were three reviewers, none of whom have requested to waive their anonymity at present, and are thus listed as A, B, and C. After initial review (1A), the author submitted a revised submission and response (1B). After a second round of review by Reviewers A and C (2A), the assigned Ledger editor determined that the author had adequately addressed the reviewer concerns. The assigned ledger editor asked the author for minor revisions, which once addressed completed the peer-review process. Author's responses in 1B are in bullet form.

1A. Review, First Round

Reviewer A:

This work analyzes a proof-of-stake algorithm addressing several questions -- how forgers are chosen (uniform vs. exponential), whether a rational forger should split itself into several smaller forgers, and what advantage it can gain by gaming the system and choosing one of its forgers such that its own forgers are chosen in future steps. The paper addresses important issues, reaching non-trivial results.

The paper's presentation does not do the content justice. Most prominently, a background section is completely missing. The paper revolves around a proof-of-stake protocol. There is no such protocol that is widely believed to be secure, despite its de-facto usage in the NXT coin. I don't see this as a deal-breaker. Good science is often done under assumptions that are later made realistic, but these assumptions must be made clear.

A comprehensive background section should cover (at least):

- Background on PoS (assumptions, algorithm, how time is treated)
- Known issues with PoS ("there is no stake", how is failure detection performed, how forks are resolved)

†Serguei Popov (popov@ime.unicamp.br) is a Professor in the Department of Statistics, Institute of Mathematics, Statistics and Scientific Computation, University of Campinas – UNICAMP, Brazil
*1FCFYiUbL3KajNvZm4W2wTSLrUimF1Lpea

Details:

Title: Capitalize

Abstract: from the probabilistic --> from a probabilistic

Intro: as an example the Nxt. --> as an example the Nxt protocol

The next Section 2 --> The next section

splitting strategies: used before definition

Footnote 1: This is an important note that should appear in the main text. Also, please detail -- what attacks are you referring to?

page 4: accessing --> assessing?
b:=b1 need not be very small: please explain more formally. Is this b1 >> bi (for i>1)?
It is elementary to show --> We observe

Figure 1: Please make BW compatible.

Page 7 (and the subsequent conclusions): There is little point in summarizing each section's information. Please leave only conclusions that were not derived before. In fact, I suggest moving all such content to a single Conclusion section at the end, as commonly done. Specifically for the conclusion in page 7 -- a missing point is that the system incentivizes participants to store their money together, leading to centralization. This is a strong point.

Section 4: find out, how --> find out how (remove comma)

Page 9: Take n=1000000: This seems like a rather high value. Is it practical? Where did you derive this value from?

Conclusion of section 4: Please explain importance, preferably with realistic numbers.

Section 5: Missing intro/segue.
being W --> W being

Page 10: does not have enough power to attack the network --> unlikely will succeed (or something of that sort)

Page 11: using the so-called... More specifically: remove, leaving simply "using Chebyshev's inequality

Page 12: for b=0.1 we have \delta(b)~=0.439 --> \delta(0.1)~=0.439 (and similarly everywhere else)

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.46

Conclusion of Section 5, bullet 2: How do they compare?

Section 6: Ban for non-forging: please explain
the attacker would have to split... anyway: do you mean without cost? "Anyway" is unclear.
the money of the attacker are distributed --> the money of the attacker is distributed
It is reasonable, however...: Please provide a backing for the claim, or simply say "we assume"

Page 14: remarkable fact --> proposition

Prop. 6.1: many accounts: please define formally

Proof of 6.1: Please be more detailed and spell out the steps of the proof

Page 14: money are --> money is

Page 15: parenthetical discussion of banning: please explain, banning was not formally explained.

Page 16: cf. e.g. [2] for the general theory: Please provide more details to make the paper self-contained.

Figure 5: axis labels missing. Please fix figure to show the curve hitting (0.5, 1).

page 17: (the author thanks...): remove
"Under current implementation": Should be under "the" current implementation". But anyway, isn't the current implementation of NXT using the uniform algorithm, which is not what was analyzed here?

Page 18, 4: This claims in this comment were not discussed before and are not backed by the paper. Moreover -- the suggested strategy would motivate a large attacker to cause the system to increase the minimum size thus removing small honest parties.
5: remove.

Paper conclusion is missing.

Reviewer B:

The results themselves are interesting and give important analysis on PoS type scheme. One major disadvantage of this paper is its structure. Author should restructure the paper according to statements of problems, taxonomy of attacking methodology and well-grained analysis structure and conclusion.

Reviewer C:

The paper is about cryptocurrency, Nxt.

Like Bitcoin, Nxt is block-chain-based cryptocurrency; it is based on P2P and is decentralized.
However, the underlying design concept of it, "Proof-of-Stake", is different from that of Bitcoin, "Proof-of-Work".

In "Proof-of-Stake" currency, the total amount of coins is fixed in advance and never changes; there is no client puzzle and no one mines a new coin.

In this type of currency, users have their own accounts and in each block generation phase, some existent account is randomly selected, and receives the intermediate fee in each transaction in this period as a reward.

This intermediate fee is the alternative incentive for generating new blocks in the "Proof-of-Stake" currency.

Note that in a jargon around the "Proof-of-Stake" community, to get the right of gaining intermediate fee is called "forging" because "forging" means the place to heat metal in order to make something (such as a coin).

In the "Proof-of-Stake" concept, the probability that an account can get the right of forging increases if the account has more coins (or in a jargon, "stake").

Ideally (for the "Proof-of-Stake" concept), the probability should be proportional to the stake of the account.

Although there were cryptocurrencies based on hybrid of both "Proof-of-Stake" and "Proof-of-Work" concepts, Nxt is the first cryptocurrency which is based only on "Proof-of-Stake" concept.

[The Subject of This Paper]

In the submitted paper, the authors study security of Nxt.

The authors mainly concern with what will happen when an attacker can gain large stake (compared to other honest users) and/or an attacker can exploit multiple accounts. (Making multiple accounts is possible without deviating from the Nxt protocol).

Let N be the number of active users of Nxt and b \in [0,1] be the "normalized" stake of the attacker, that is, (the stake of the attacker)/(the sum of stakes of all active users).

[Sec 2 and 3]

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.46

(result)

In Section 2 and 3, the authors consider the case where b is very large compared to the normalized stakes of the other users and show that the probability P that the attacker can get right of forging is b + b^2 +O(b^3), when N -> \infty and b\to 0.

As mentioned before, the above probability in an ideal "Proof-of-Stake" cryptocurrency should equal to b.

Hence, the probability, P=b+b^2+O(b^3), for Nxt is larger than the ideal one.

(technical comments)

The authors should exemplify the case when (N,b) becomes (\infty,0), and why this is a valid case to consider.

The authors should clarify what are the consequences of the result.

I agree that P=b+b^2+O(b^3) > b.

But I cannot understand the authors want to say whether "Nxt is good because P is almost same as b" or "Nxt is not good because P is much larger than b".

[Sec 3.1]

(result)

The authors consider the case where an attacker has multiple accounts.

Then they show that the multiple accounts are meaningless for gaining right of forging because the probability P that the attacker can get right of forging is largest when the attacker deposits all of its stake in one account.

(comment)

This is good result because the result implies that we can concentrate on studying security under the condition that an adversary has only one account. From organization of the paper, I wonder why this is a subsection.

[Sec 4 and 5]

(result)

The authors estimate the expected length that a user forges consecutive blocks of chain. Then the authors estimate the probability Q that an attacker succeeds in forging m consecutive blocks of chain offline.

Then they show that

- if b < 1/3, Q exponentially converges to 0, when the number m of blocks becomes \infty,
- if b \ge 1/3, Q is almost 1 even when m is large.

(Technical comment)

This is good analysis too, but again, I cannot understand what is the consequence of the analysis.

The authors say that m=10 or so (in page 10).

This means that even if b=0.05, Q becomes \delta(0.05)^{m} =0.2^{10} = 1/10^7.

1/10^7 is not very small when we compare it with, e.g., the probability 1/10^24 that the attacker succeeds in breaking a 160-bit hashing.

Security of Nxt depends on both the significance of the incident (succeeding in counterfeiting blocks) and the probability Q that the incident happens.

Hence, the author has to clarify what damage would happen if an attacker succeeds in counterfeiting blocks of chains and has to give the conclusion whether 1/10^7 is large or not.

Moreover, the validity on the assumption of the size of m(=10).

[Sec 6]

I failed to understand the details of this section.

Probably, the authors want to say that when an attacker succeeds in taking right of forging (e.g. by depositing all of its stake in one account as described in Sec 3.1), the attacker can maximize the probability that it takes right of forging of the next phase by splitting its stake into lots of accounts.

In particular, if an attacker succeeds in getting 50% of all existent stakes, it can forge all the blocks.

However, this conflicts with the result described in Section 3.1. It was difficult to understand if there is any assumption difference.

[Other Results]

The authors consider another algorithm, called Exp-algorithm, to determine who can have right of forging.
Exp-algorithm is different from the algorithm of Nxt, called U-algorithm, to determine it.

The authors show that Exp-algorithm is better than U-algorithm in some ways.

Again, I cannot understand what is the consequence of this claim.

Do the authors want to say that we should create new cryptocurrency based on Exp-algorithm?

[Editorial Comments]

The paper is not self-contained.

The paper should be understandable even for readers who are not familiar with cryptocurrency.

The author has to explain technical procedures, especially the details of "Proof-of-Stake" and the details of the block chain algorithm of Nxt.

Perhaps many of my comments are due to lack of this explanation.

[Evaluation]

The result is interesting and I think that the paper is acceptable.

However, as mentioned before, the authors do not write the consequences of their results in using Nxt.

Due to the lack of them, I do not recommend accepting the paper as it is.

1B. Author's Response

I've rewritten the introduction (it became twice as big) to address several referee's comments. In particular, the changes to in the introduction include:

- a basic comparison of the PoW and PoS protocols, and some notes on the terminology used
- added an observation that the subject of the paper is the so-called "pure PoS", and added a note on other PoS versions (such as those that include coin-age)
- added a note on some common attacks on the PoS-protocol; however, discussing all these attacks is well beyond the scope of this paper. Added references to some papers of The Consensus Research, where these attacks are discussed at length.
- also, added some explanations on the model (in particular, about the time) to Section 2

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.46

Reviewer A:

Title: Capitalize

- left it as it was, for now, since in any case it's not yet clear which layout The Ledger will use

Abstract: from the probabilistic --> from a probabilistic

- corrected

Intro: as an example the Nxt. --> as an example the Nxt protocol

- corrected

The next Section 2 --> The next section

- corrected

splitting strategies: used before definition

- rewritten this part

Footnote 1: This is an important note that should appear in the main text. Also, please detail -- what attacks are you referring to?

- moved the footnote to the main text, and added some explanations on the attack to (the beginning of) Section 4

page 4: accessing --> assessing?

- corrected

b:=b1 need not be very small: please explain more formally. Is this b1 >> bi (for i>1)?

- rewritten this part, to make it more clear

It is elementary to show --> We observe

- corrected

Figure 1: Please make BW compatible.

- I've remade all the figures in the paper

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.46

Page 7 (and the subsequent conclusions): There is little point in summarizing each section's information...

- I've removed all conclusions after the sections, and wrote the final conclusion instead

Section 4: find out, how --> find out how (remove comma)

- corrected

Page 9: Take n=1000000: This seems like a rather high value. Is it practical? Where did you derive this value from?

- explained that this is, in fact, close to the current size of the Nxt's blockchain

Section 5: Missing intro/segue.

- added a couple of paragraphs in the beginning, to explain what us going on

being W --> W being

- corrected

Page 10: does not have enough power to attack the network --> unlikely will succeed (or something of that sort)

- corrected

Page 11: using the so-called... More specifically: remove, leaving simply "using Chebyshev's inequality

- corrected

Page 12: for b=0.1 we have \delta(b)~=0.439 --> \delta(0.1)~=0.439 (and similarly everywhere else)

- corrected

Conclusion of Section 5, bullet 2: How do they compare?

- they are now plotted together on Figure 3

Section 6: Ban for non-forging: please explain

- explained (2nd paragraph of Section 6)

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.46

the attacker would have to split... anyway: do you mean without cost? "Anyway" is unclear.

- rewritten this part

the money of the attacker are distributed --> the money of the attacker is distributed

- corrected

It is reasonable, however...: Please provide a backing for the claim, or simply say "we assume"

- put "we assume"

Page 14: remarkable fact --> proposition

- corrected (with "result" instead of "proposition")

Prop. 6.1: many accounts: please define formally

- introduced the notion of "splitting the account to dust" just before Proposition 6.1

Proof of 6.1: Please be more detailed and spell out the steps of the proof

- added some details just before the first display on p. 15

page 14: money are --> money is

- corrected

Page 15: parenthetical discussion of banning: please explain, banning was not formally explained.

- put an explanation earlier (2nd paragraph on p. 14)

Page 16: cf. e.g. [2] for the general theory: Please provide more details to make the paper self-contained.

- wrote a small intro to Galton-Watson branching processes (1st paragraph on p. 17)

page 17: (the author thanks...): remove

- removed

"Under current implementation": Should be under "the" current implementation". But anyway, isn't the current implementation of NXT using the uniform algorithm, which is not what was analyzed here?

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.46

Page 18, 4: This claims in this comment were not discussed before and are not backed by the paper. Moreover -- the suggested strategy would motivate a large attacker to cause the system to increase the minimum size thus removing small honest parties.

5: remove.

- removed

Paper conclusion is missing.

- not anymore

Reviewer C:

The authors should exemplify the case when (N,b) becomes (\infty,0), and why this is a valid case to consider

- (N,b) doesn't really become (\infty,0), we only consider the limit as b\to 0.

The authors should clarify what are the consequences of the result. I agree that P=b+b^2+O(b^3) > b. But I cannot understand the authors want to say whether "Nxt is good because P is almost same as b" or "Nxt is not good because P is much larger than b".

- added explanations at the bottom of p. 5

Sec 3.1
From organization of the paper, I wonder why this is a subsection.

- this is because it discusses the effect of splitting on the probability of generating a block (which is the subject of Section 3)

Sec 4 and 5
(Technical comment)
This is good analysis too, but again, I cannot understand what is the consequence of the analysis. The authors say that m=10 or so (in page 10). This means that even if b=0.05, Q becomes \delta(0.05)^{m} =0.2^{10} = 1/10^7. 1/10^7 is not very small when we compare it with, e.g., the probability 1/10^24 that the attacker succeeds in breaking a 160-bit hashing.

Security of Nxt depends on both the significance of the incident (succeeding in counterfeiting blocks) and the probability Q that the incident happens. Hence, the author has to clarify what damage would happen if an attacker succeeds in counterfeiting blocks of chains and has to give the conclusion whether 1/10^7 is large or not. Moreover, the validity on the assumption of the size of m(=10).

- m=10 was a sort of "rule of thumb" in the Nxt for accepting transactions (similar to "wait 6 blocks" in Bitcoin). It is true that 1/10^7 is not very small compared with the probability that the attacker succeeds in breaking a 160-bit hashing, but the consequences of a single double-spend (cf. the beginning of Section 4) are not so dramatic either. Besides, if the stakes are high, the merchant may prefer to wait for more confirmations before accepting the transaction for good.

Sec 6
I failed to understand the details of this section. Probably, the authors want to say that when an attacker succeeds in taking right of forging (e.g. by depositing all of its stake in one account as described in Sec 3.1), the attacker can maximize the probability that it takes right of forging of the next phase by splitting its stake into lots of accounts.

- The idea of this attack is that the attacker tries to take _several_ places in front of the forging queue. Since the forging algorithm works in a pseudo-random way, the attacker is able to predict the next forging queue as a function of the account he _chooses_ to forge the current block (since there are several accounts of his in front of the queue, he has this choice). So, he can take the "best" (for him) case, when there are some accounts of his in front of the next forging queue as well. We argue then that, if the stake of the attacker is high enough, he can eventually repeat all this ad infinitum. I wrote some more clarifications on this in the text.

However, this conflicts with the result described in Section 3.1. It was difficult to understand if there is any assumption difference.

- In Section 3.1 we discussed only the probability of generating the next block, in the simpler situation when the entity that forges the current block doesn't play any "forging games" as the one described above.

Other Results

The authors consider another algorithm, called Exp-algorithm, to determine who can have right of forging. Exp-algorithm is different from the algorithm of Nxt, called U-algorithm, to determine it. The authors show that Exp-algorithm is better than U-algorithm in some ways. Again, I cannot understand what is the consequence of this claim. Do the authors want to say that we should create new cryptocurrency based on Exp-algorithm?

- As noted on the bottom of p. 5, the U-algorithm is acceptable as well, since its implementation is simpler.

2A. Review, Second Round

Reviewer A:

The author has responded to my concerns. I would recommend however that it goes through proof reading with an English speaker before publication for correctness and style.

I do not need to see another version, but please pass the following two comments to the author:

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.46

1. "Also, all the coins are created in the beginning": Is this a necessary property of a PoS protocol? Why?

2. Conclusion: at this point the reader has already read the paper. Use it to summarize the actual results and offer insights -- big picture, not just mentioned what you have studied, as you can do in the intro.

Reviewer C:

The paper provides a probabilistic analysis of pure-PoS used in Nxt.

The analysis is on:

1) The probability that an account can forge a block.

2) The probability of forging a consecutive blocks. Formula is given, but no specific consequences are provided.

3) The probability of genenrating a longer sidechain.

4) Effect of account splitting

The paper provides good analysis, so the paper should be published, given the following modifications:

1. Summarize, in the Conclusion section, the observations provided through the analysis, especially that from Section 3.1 and Section 5.

2. Regarding the formula (8) provided in Section 4, provide more numeric examples based on different x and p and discuss their consequences.

3. In the last sentence of Section 5, please check if inequality symbols are correct.

Ledger is published by the University Library System of the University of Pittsburgh as part of its D-Scribe Digital Publishing Program and is cosponsored by the University of Pittsburgh Press.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.46