LEDGER
ledgerjournal.org

ISSN 2379-5980 (online)
associated article DOI 10.5915/LEDGER.2020.195

Real-Time Block Rate Targeting: Open Review

Thomas M. Harding[*][†]

Reviewers: Reviewer A, Reviewer B

Abstract. The final version of the paper "Real-Time Block Rate Targeting" can be found in Ledger Vol. 5 (2020) 11-19, DOI 10.5915/LEDGER.2020.195. There were two reviewers involved in the review process, neither of whom have requested to waive their anonymity at present, and are thus listed as A and B. After initial review by Reviewers A and B, the submission was returned to the authors with feedback for revision (1A). The authors responded (1B) and resubmitted their work. It was once again sent to Reviewers A and B, who indicated that the revisions made were sufficient to address their concerns, thus ending the peer review process. Author responses are bulleted for clarity.

1A. Review

Reviewer A

Does this paper represent a novel contribution to cryptocurrency or blockchain scholarship?

Yes

If you answered "yes" to the previous question, in one sentence, describe in your own words the novel contribution made by this paper:

This is probably the first peer-reviewed paper describing how to adjust targets based on the block's own solvetime.

Is the research framed within its scholarly context and does the paper cite appropriate prior works?

Yes

[*] 1BaKh3EaAjGoaP8BH9AZ4jG7VD3HaC3xoG.
[†] T. M. Harding (tom@chain2.org) is an independent researcher and contributor to open-source software development projects.

Please assess the article's level of academic rigor:

Excellent (terms are well defined, proofs/derivations are included for theoretical work, statistical tests are included for empirical studies, etc.)

Please assess the article's quality of presentation:

Excellent (the motivation for the work is clear, the prose is fluid and correct grammar is used, the main ideas are communicated concisely, and highly-technical details are relegated to appendixes).

How does the quality of this paper compare to other papers in this field?

Top 5%

Please provide your free-form review for the author in this section:

Is it correct to say tn is a random variable?

Hashrate will often drop to 1/2 if there is a 10% drop in the price/difficulty ratio, not linear as stated.

It would be nice to mention in last paragraph of page 2 the recent BCH oscillations (which are a great example of 600% changes in hashrate resulting from 10% drops in price/difficulty ratio). As is, it sort of indicates BCH's algo is merely lacking optimization.

I suspect changing the un / T in eq 14 to e^(-un) / e^(-T) will make section 3.3 and appendix 2 unnecessary, greatly simplifying the most difficult part of the paper to understand. This maps asymmetrical time back to a linear random variable so that repeated multiplications of this ratio in re-targeting do not result in that error, making the algorithm more precise at each change.

I have two major complaints about the article, both of which are difficult to address. The 1st is that it motivates very large changes in hashrate during the block that will change lambda(t) because it is also a function of hashrate, which changes the derivations. Average solvetime adjustment that is determined by the CDF may remain accurate despite hashrate changes causing the PDF to be much much tighter with a sharper peak around T. I am concerned a tight peak will significantly increase the orphan rate. I would not be surprised if k=2 is the largest value that should be used. It's a difficult thing to address because we do not know how quickly hashrate can change coins without losses.

The other issue is that timestamp manipulation is not addressed. This will be the biggest concern of most readers. A miner can starting mining a new block with a timestamp set in the future such as tn=T so the difficulty is lower. If he finds it early, he has to wait before releasing it to nodes with good clocks who will reject it until T arrives. He risks losing the block to honest miners, but he is motivated to start mining the next block so he may still orphan an honest block that beat him on the first block. So a miner with > 50% HR is much

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2020.195

more strongly motivated to do a selfish mine (large miners working on small coins usually choose not to selfish mine). Competition for cheating may decrease profits such that normal users may do no cheating and lose very little, but that's only a hope. If a profitable scheme is devised, all miners could be advised to follow the scheme, at which point a smart cheater would devise a new scheme, possibly cycling between schemes. The complexity may increased orphan rate and selfish mining even more, probably reducing total hashrate which reduces the security of the coin.

Reviewer B

Does this paper represent a novel contribution to cryptocurrency or blockchain scholarship?

Yes

If you answered "yes" to the previous question, in one sentence, describe in your own words the novel contribution made by this paper:

Introduces a DAA that adjusts every second to increase likelihood of finding a block the longer one hasn't been found..

Is the research framed within its scholarly context and does the paper cite appropriate prior works?

Important references are missing

Please assess the article's level of academic rigor:

Excellent (terms are well defined, proofs/derivations are included for theoretical work, statistical tests are included for empirical studies, etc.)

Please assess the article's quality of presentation:

Excellent (the motivation for the work is clear, the prose is fluid and correct grammar is used, the main ideas are communicated concisely, and highly-technical details are relegated to appendixes).

How does the quality of this paper compare to other papers in this field?

Top 20%

Please provide your free-form review for the author in this section:

This is an interesting, well-written paper. To summarize: the proposed DAA adjusts block target difficulty every second to make finding a block at the network level easier the longer it has been since the last block has been found. The DAA can be modified with the parameters of a, k, and T to increase the probability that the new block will be found after T seconds. In addition to the block target (g(n)), when the block is found, its un-adjusted target difficulty is calculated (G(n)). This un-adjusted target difficulty of block n is used in the target calculation of block n+1. In this way, inter-block times are pushed towards the target block time with g(n), but the difficulty also can respond to hash fluctuations at every block with the G(n) calculation.

A few comments, suggestions, and questions I have (in no particular order) are:

1) I don't understand how the 2.2% and >5% numbers are calculated or where they are drawn from in this sentence:
"That decision reduced the coefficient of variation from 100% to 2.2%, but has resulted in a >5% overshoot versus target of the long term block production rate in the environment of ever-increasing hashrate."

2) Some related work that would be good to reference in the final version are:
- A Lucas Critique to the Difficulty Adjustment Algorithm of the Bitcoin System
- Difficulty control for blockchain-based consensus systems
- The upcoming FC'20 paper, Selfish Mining Re-examined, since it examines various current DAAs.

3) In section 4, you state that the block comparator is changed to reference the chainwork of the parent and that this prevents simple orphaning attacks. I don't think this is correct because at the first sign of a fork, how do miners choose which block to mine on? Both blocks share the same parent so choosing the block with higher parent chainwork here doesn't make sense. However, miners will want to mine on the older block because your algorithm makes mining on the older block easier since the target will necessarily be lower than mining on the newer block. The orphaning attack of starting a competing chain with a block with an earlier timestamp seems feasible to me, contrary to what you claim in this section.

4) Depending on your page limits, I think you should go more into how your DAA fares against deviant mining behavior, including fleshing out section 4 on forks and competing chains. For example, selfish mining seems likely with you DAA. If a selfish miner finds a block before the honest miner and does not publish it until the honest miner finds a block, then the honest miners will choose to mine on the selfish chain since that block has an earlier timestamp and so the target on this chain is lower than on the honest chain. This is similar to the phenomenon observed in Selfish Mining Re-examined with ETH, since the ETH DAA also modifies difficulty at every block.

5) For the graphs, I would like to see axis labels and equation lines with different colors and styles for color blindness.

1B. Author Responses

Reviewer A

Is it correct to say tn is a random variable?

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2020.195

- Yes, t_n is a measurement whose actual value can't be predicted with certainty, but which may be reasoned about if its distribution is known or supposed.

Hashrate will often drop to 1/2 if there is a 10% drop in the price/difficulty ratio, not linear as stated.

- Thank you, this inaccurate and unnecessarily specific quantification has been removed.

It would be nice to mention in last paragraph of page 2 the recent BCH oscillations (which are a great example of 600% changes in hashrate resulting from 10% drops in price/difficulty ratio). As is, it sort of indicates BCH's algo is merely lacking optimization.

- Thank you. Section 1.2 now cites academic results that document the BCH oscillations from a theoretical standpoint.

I suspect changing the un / T in eq 14 to e^(-un) / e^(-T) will make section 3.3 and appendix 2 unnecessary, greatly simplifying the most difficult part of the paper to understand. This maps asymmetrical time back to a linear random variable so that repeated multiplications of this ratio in re-targeting do not result in that error, making the algorithm more precise at each change.

- It's not repeated multiplications that result in the need to adjust T. As Rosenfeld found, the expected time of the very next block after a difficulty adjustment is different from the intended adjustment -- by the full adjustment magnitude -- because difficulty adjustment involves division by a random variable. I don't quite follow your exponential ratio suggestion; "mapping back asymmetrical time" is what we do in section 3.2, and any modification to the results would have to start with a motivation and a formulation.

I have two major complaints about the article, both of which are difficult to address. The 1st is that it motivates very large changes in hashrate during the block that will change lambda(t) because it is also a function of hashrate, which changes the derivations. Average solvetime adjustment that is determined by the CDF may remain accurate despite hashrate changes causing the PDF to be much much tighter with a sharper peak around T. I am concerned a tight peak will significantly increase the orphan rate. I would not be surprised if k=2 is the largest value that should be used. It's a difficult thing to address because we do not know how quickly hashrate can change coins without losses.

- I agree that miners might choose to join each block at some positive second which they deem the profitable point of entry, rather than at second zero. I also agree that this behavior leads to a tightening of the distribution, since their foregonesince block production in earlier seconds is introduced at a time of lower difficulty. Further, I agree that this will not throw off the targeted mean block time -- any more or less hashrate applied at any given second still has its effect on block production modulated

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2020.195

by the correct difficulty. And finally, I agree that the narrower distribution of block times inevitably means a greater percentage of stale blocks.

- To explore the stale block probability, albeit only at constant hashrate, I considered a number of equally-sized miners, which is a poor-case scenario for that number of miners. Then I considered propagation times ranging from 1s to 5s. With all the propagation technologies (compact blocks, Graphene, XThin) and networks (Fibre, Falcon, BloXRoute) I expect propagation to be between 1 to 2 seconds these days even for large (by today's standsards) blocks.

- A stale block is one that is found within <time tolerance> seconds after the first miner to find the block. It's also possible for multiple miners to find a stale block for the same first-block - this is accounted for.

- As expected, RTT has significantly more stale blocks because of the narrower distribution.

- Each of these numbers represents 10,000 trials and this runs in less than a minute. The numbers don't change much after 20 miners because the expected time is so long for each miner.

- RTT (k=6) Stale Block Probability

k=6, Miners / Time Tolerance	1	2	3	4	5
2	0.0046	0.0094	0.0111	0.0181	0.0201
4	0.0077	0.0131	0.0197	0.0249	0.0342
6	0.007	0.0135	0.0201	0.0283	0.0384
8	0.0052	0.0167	0.026	0.0324	0.036
10	0.0078	0.0144	0.0237	0.0321	0.0379
12	0.0084	0.0163	0.0227	0.0332	0.0382
14	0.0072	0.0156	0.0262	0.0296	0.041
16	0.0075	0.0184	0.0222	0.0317	0.0414
18	0.0077	0.0178	0.0236	0.0314	0.0441
20	0.0086	0.0161	0.021	0.0325	0.0441

- For comparison,
- Bitcoin (k=1) Stale Block Probability

k=1, Miners / Time Tolerance	1s	2s	3s	4s	5s
2	0.0009	0.0019	0.0023	0.0024	0.0049
4	0.0012	0.0028	0.0036	0.0058	0.0058
6	0.0011	0.0028	0.0046	0.0046	0.0064
8	0.0011	0.0023	0.005	0.0057	0.0068
10	0.0015	0.0029	0.0045	0.0059	0.0074
12	0.001	0.0023	0.005	0.0049	0.0076
14	0.0009	0.0045	0.0041	0.0074	0.0071
16	0.0011	0.0042	0.0043	0.0057	0.0074
18	0.0014	0.003	0.0046	0.0059	0.0069
20	0.0018	0.0025	0.004	0.0067	0.0075

The other issue is that timestamp manipulation is not addressed. This will be the biggest concern of most readers. A miner can starting mining a new block with a timestamp set in the future such as tn=T so the difficulty is lower. If he finds it early, he has to wait before releasing it to nodes with good clocks who will reject it until T arrives. He risks losing the block to honest miners, but he is motivated to start mining the next block so he may still orphan an honest block that beat him on the first block. So a miner with > 50% HR is much more strongly motivated to do a selfish mine (large miners working on small coins usually choose not to selfish mine). Competition for cheating may decrease profits such that normal users may do no cheating and lose very little, but that's only a hope. If a profitable scheme is devised, all miners could be advised to follow the scheme, at which point a smart cheater would devise a new scheme, possibly cycling between schemes. The complexity may increased orphan rate and selfish mining even more, probably reducing total hashrate which reduces the security of the coin.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2020.195

- I've worked with another researcher to quantify this and explore it a bit, and there may be new research specifically on the idea of "defacto future mining" which you describe. It seems to be mitigated by the attractiveness of an even simpler strategy, which is to stay away from the RTT blockchain during the early part of the block search, until that moment when it is more profitable to mine than the competition.

- I hasten to point out that defacto future mining is not timestamp manipulation, because it does not cause blockchain time to differ from wall clock time. There is every reason to think that RTT blockchain time will be far more resistant to that kind of timestamp manipulation that other blockchains seen to date.

- The complexity you refer to is something that may require experience with a real RTT blockchain, to gain insight and direction.

Reviewer B

1) I don't understand how the 2.2% and >5% numbers are calculated or where they are drawn from in this sentence:

"That decision reduced the coefficient of variation from 100% to 2.2%, but has resulted in a >5% overshoot versus target of the long term block production rate in the environment of ever-increasing hashrate."

- The coefficient of variation (standard deviation divided by mean) of an exponential random variable is 100%. For the sum of N such random variables, it is 1/sqrt(N), which for N=2016 is 2.2%.

- The 5% block rate overshoot figure was measured at BTC block height 596230, whose timestamp would have appeared on a block about 5% lower if an average rate of only 1 block per 10 minutes had been maintained since genesis.

2) Some related work that would be good to reference in the final version are:

- A Lucas Critique to the Difficulty Adjustment Algorithm of the Bitcoin System
- Difficulty control for blockchain-based consensus systems
- The upcoming FC'20 paper, Selfish Mining Re-examined, since it examines various current DAAs.

- Thank you for the excellent references.

- I had read a draft of the first article "A Lucas Critique..." and failed to recall it when doing this work. I'm sure it influenced me. That alone is sufficient reason to add it as a reference, but in addition, on re-reading it, I have made revisions to section 1.2 and cited this article therein.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2020.195

- I also found the second article "Difficulty Control..." interesting and relevant, but chose not to reference it, since it focuses on a particular pattern of hashrate evolution and I had not read it when doing this work.

- Unfortunately, the third paper was not yet available for me to review.

3) In section 4, you state that the block comparator is changed to reference the chainwork of the parent and that this prevents simple orphaning attacks. I don't think this is correct because at the first sign of a fork, how do miners choose which block to mine on? Both blocks share the same parent so choosing the block with higher parent chainwork here doesn't make sense. However, miners will want to mine on the older block because your algorithm makes mining on the older block easier since the target will necessarily be lower than mining on the newer block. The orphaning attack of starting a competing chain with a block with an earlier timestamp seems feasible to me, contrary to what you claim in this section.

- Thank you. The language in section 4 ("denies the opportunity") was too strong. What I am confident of is that system is more resistant to reorganization attack with the parent chainwork rule, that it would be without that rule. I have updated the language.

- I agree that attack is feasible, and more so than in bitcoin. One consequence is, other things equal, a given number of confirmations under RTT may confer the security of one fewer confirmation under bitcoin.

4) Depending on your page limits, I think you should go more into how your DAA fares against deviant mining behavior, including fleshing out section 4 on forks and competing chains. For example, selfish mining seems likely with you DAA. If a selfish miner finds a block before the honest miner and does not publish it until the honest miner finds a block, then the honest miners will choose to mine on the selfish chain since that block has an earlier timestamp and so the target on this chain is lower than on the honest chain. This is similar to the phenomenon observed in Selfish Mining Re-examined with ETH, since the ETH DAA also modifies difficulty at every block.

- While this additional research is essential to incorporating RTT into useful systems, it is a large undertaking and I had to leave it out of scope for this article.

5) For the graphs, I would like to see axis labels and equation lines with different colors and styles for color blindness.

- The graphs have been made larger and higher-contrast. I hope they are easier to read for everyone.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2020.195

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.2020.195