LEDGER
ledgerjournal.org

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

Subchains: A Technique to Scale Bitcoin and Improve the User Experience: Open Review

Author: Peter Rizun*†

Reviewers: Reviewer A, Reviewer B, Reviewer C

Abstract. The final version of the paper "Subchains: A Technique to Scale Bitcoin and Improve the User Experience" can be found in Ledger Vol. 1 (2016) 38-52, DOI 10.5915/LEDGER.2016.40. There were three reviewers who responded, 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 responses (1B). The revised submission was reviewed by the assigned Ledger editor who determined that the author had adequately and substantively addressed the reviewers' concerns, thus completing the peer-review process. Author's responses in are in bullet form.

1A. Review, Initial Round

Reviewer A:

Thank you for this paper.

I've not been able to comment on the math, due to my lack of math-skills. But I do understand the reasoning and conclutions. I hope that's fine.

The main purpose of the paper is to reduce orphaning risks and to make 0-conf double spends less probable to succeed. The Orphaning risk is reduced by minimizing the amount of data sent at block propagation time. The 0-conf security is increased by miners' incentives to stay on the subchain with the most fees.

The idea of weak blocks is very interesting. The obvious benefit from subchains (and other weak block proposals) is reduced orphaning risks. But the less obvious (that I hadn't though of before) effect is that 0-conf transactions may be safer, but I think the paper needs to be more convincing on that matter since the paper lacks elaboration on different types of double spend attacks.

Major comments:

† P. R. Rizun, Ph.D. (peter.rizun@gmail.com) is Chief Scientist for Bitcoin Unlimited and resides in Vancouver, Canada.
*1BWZe6XkGLcf6DWC3TFXiEtZmcyAoNq5BW

In section 1, one of the assumptions made is that "Miners are rational, short-term profit-maximizing agents". I'm missing some elaboration on how probable this is to be true, both today and in the future. Today, many miners use default settings of bitcoind, meaning they do what the developers considers best for the network as a whole, but when margins tightens, miners might be forced into tweaking their policies and possibly make agreements with each other and with certain payment processors, merchants, exchanges and so on.

I think the incentives for miners to propagate its weak blocks are unclear and needs more elaboration. For example, would it really be more profitable for me to publish a weak block if I find one? If I find the next strong block, I would benefit from it only if it's immediately after my weak block. If there's another weak block after my block, it made no difference at all. On the other hand, if I don't find the next strong block, I've just given away fast block propagation, ie low orphan risk, and the fee benefit of weak blocks, to someone else.

Section 8: This section assumes that an attacker will have to (the word "must" in the first sentence of the section) increase its weak block size, by adding transactions (for the fees), faster than the rest of the network in order to produce a higher-fee subchain. But the main goal for the attacker is to perform a successful double spend. Assuming all miners (but the attacker) will build off the highest-fee subchain, all the attacker have to do is create a weak block with ONLY the double spend, but with a higher fee than the highest-fee subchain. If he succeeds, all the hashing power will switch over to his double spend subchain. If he fails, and the next weak block is built off of the honest chain, the attacker will simply create a new block template with a new double spend with the fee increased to be bigger than the new highest-fee subchain. In a sense, double spending becomes easier, because the attacker have several chances to succeed: one per weak block produced in the honest subchain. To minimize the fee of the double spend, the attacker can of course add as many other transactions he wishes as long as they don't make the block too big (with respect to orphaning risk) or conflict with his subchain.

Section 8: If my previous comment is valid, then much of section 8 (the statement that it helps 0-confiramtion transaction security) is invalid for high value double spends. For low value double spends it might still work as described in the paper, because there is no sense in setting the fee higher than the value of the payment. The calculus must take into consideration all reasonable types of double spend attacks.

The assumption (1) in section 1 implies that miners accept replace-by-fee transactions. This means that a transaction can be cheaply replaced (just as today) until it's included in a weak block of the highest-fee subchain. When it's included in the highest-fee subchain the cost of replacing it depends on how deep (fee-wise) in the subchain it is buried, and if that cost increase is bigger than the fee increase of the replacement transaction, it's not worth it.

It seems like RBF is incompatible with subchains since you can only add transactions to a subchain, not remove (or swap) them. In order to replace an opt-in replacable transaction, you must either build a higher-fee subchain or not include replacable transactions at all in the subchain, meaning they are put only in strong blocks. You cannot foresee if your block

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

template is going to be a strong or weak block, so you cannot ever include a replacable transaction without making it unreplacable. Another way of dealing with replacement would be to simply add the replacement transaction in a subsequent weak block of the subchain and let that mean that it should replace the previous transaction. But this would of course make double spending just as hard/easy as it is today, so the argument about increasing 0-conf security becomes void.

Minor comments:

- The "List of Symbols" section is very helpful.
- Section 4, first paragraph: "A miner is thus financially incetivized to build off the highest-fee subchain": Maybe not true, a miner may have other transactions it wants to mine that actually conflicts with the highest-fee subchain. It might be that a miner have an agreement with a payment provider or other organization that forces them to select transactions in other ways. The next sentence then falls too "Since all miners have the same incentives...".
- Fig 3: The picture needs a) and b), because they are referenced from text.

Reviewer B:

This paper deals with the idea of subchains. Subchains are composed of weak blocks. Weak blocks are solved by miners in the same way (strong) blocks are solved, only with a weaker (larger) target. The advantage of using subchains is that miners can progressively send smaller pieces of information to the network when building weak blocks. Hence, this has an important impact on the risk of having a weak block orphaned and hence lowers the cost of solving blocks, allowing more transactions to be included in strong blocks. The author explains the mechanism of subchains very clearly and incentives faced by miners when using it (Section 3-4-5), formally shows their advantage in terms of orphan risk and hence in terms of modifying the block space offer function (Section 6), formally computes the cost of trying to outrace the network for an attacker who would like to double-spend some coins (Section 8) and finally notes the possibility to extend this idea of subchains by nesting them (Section 9).

This paper is very interesting and of great interest for the crypto-currency community. I recommend that Ledger publishes this article. I still have a few comments that I would like the author to address.

Major remarks:

One of the major advantages of subchains that the author states in the introduction and the in conclusion is that it does not require a fork (soft or hard) in order to be implemented. It is true that, as it is presented, it only requires participating miners to agree on this protocol to understand each other. And as noted "it does require participation from a significant fraction of the network hash power in order to be useful." The study is then carried on assuming the whole set of miners agree on the subchain design. However, none of the transition (from a situation in which only a small fraction miners adopt subchains) period is discussed. I am not sure this period is only a period of unusefulness of subchains as suggested in the above quoted

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

sentence. Instead, I suspect that there are negative incentives to join the subchain movement. For instance, if only a small fraction of the miners agree on the subchain design (I am aware it is not the question the author adresses), a) either they have to broadcast at some point their strong blocks to the rest of the network or b) they don't. In case a), the fact that strong blocks built with subchains are larger than "regular" strong blocks becomes a disadvantage for the miners having adopted the subchain design ("regular" strong blocks miners are supposed to have the optimal number of transactions in their block). Of course, this would mean that this effect should be anticipated and the consequence are not clear (and certainly depends on who has an incentive to broadcast the strong block). In case b), a fork actually occurs. It is not a fork in the sense of a software change but it is still a fork with two incompatible blockchains. Then, at best, there is a coordination game between miners. This is a qualitative and certainly not complete reasoning but I would suggest the author addresses this question maybe by discussing it in the conclusion a little bit more extensively and clearify the point of a mere unusefulness of subchains when the participation of only a small fraction of the hash rate is "voting" for subchains (maybe only to say that it is not clear if the transition is feasible).

In the whole paper, the author states that one of the main advantages of subchains is to increase the number of transactions processed by Bitcoin. (eg p1: "Unlike Visa, Bitcoin’s transactional capacity is limited due to miners’ hesitation to produce blocks containing large volumes of new transactions.") In the formal study, it is transparently stated as an assumption that "The free-market equilibrium block size is smaller than the protocol-enforced block size limit (if such a limit exists)." However, today, the problem of limited number of transactions is more a problem of protocol-enforced block size limit than a orphaning risk driven fee problem. Maybe the "BLOCKSIZE LIMIT DEBATE WORKING PAPER" header does not help. In my opinion, this paper is more a proposition for a cost-cutting production function (when the blocksize limit is not binding) rather than bringing arguments in the blocksize debate. This should be more clear in this article as I guess today's context is very present in today's readers' minds (even though I understand this blocksize limit might be only a temporary issue for Bitcoin).

Minor remarks:

p3: Fig1 is quite misleading because it lets the reader think that the number of weak blocks per strong block is fixed whereas it is not the case. I would just remove this figure as the text is explanatory enough in my opinion.

p7: I am not sure about the τ_0 term in $P_{orphan}=1-e^{\frac{\tau_0}{T}}e^{zQ\DeltaT}{T^2}$ equation. Indeed, if all miners need τ_0 to spread their solution to the majority of other miners(assuming that all other miners mine empty blocks) weak or strong, then, the probability to be considered should be the one that no other miner finds a block between 0 and $\tau - \tau_0$. Maybe Figure 5 should be modified accordingly. Otherwise it means that other miners have instant spreading of their block.

p7: it should be made clear that another assumption is used here: All miners assume that other miners mine empty blocks. This has important implications in game theory (see Houy's "The Bitcoin Mining Game").

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

Fig 5: The considered size per transaction is not specified in the text. It should be for the reader to understand the link between the upper and lower scales.

p8: "the rate at which orphaning risk, M, is incurred". It is not clear here what the definitions are (for those who did not read the "market fee exists" paper by the same author). One more line defining M would certainly be welcome. (maybe even in the beginning of Section 4).

Fig8: same as Figure1: it is too "symmetrical" for not being misleading. Again, the text is explanatory enough in my opinion.

footnote 24: "Since ?supply = zRT" should be "Since ?supply = zRT^{-1}".

Equation 10: If I am not wrong, as it is proved, Equation 10 is obtained by integrating until infinity an expression that is a power series about 0. Maybe table 1 could just be obtained by means of numerical computation rather than trying to have an explicit approximate expression (as said above, maybe not that approximate as it is).

Reviewer C:

The paper is generally nicely written and clearly a great deal of work has gone into presentation, figures, etc. Still, I found several problems with the analysis in the paper.

The analysis assumes all miners have similar mempools, and are working on the same block. This may be true when attackers are not present, but what if the attacker sends many conflicting transactions to different nodes in an attempt to increase variations between mempools? How does the protocol react then?

Why would the first delta block be small, and not the size of an entire block? Miners should usually have something in the mempool that they can start hashing that would be more profitable than a small delta block.

There seems to be no accounting for the arrival of transactions with varying fees. If a high paying transaction suddenly appears in a delta block, miners may want to include it, but also to evict previously included low-fee transactions that had been in the previous version of the block being processed.

The security analysis relies on establishing that an attacker will suffer a cost from attacking. This is highly problematic, as it is already well established that double-spending attacks in Bitcoin are profitable for attackers. Any miner engaging in selfish mining occasionally creates long secret chains that can be used to double spend. This strategy is profitable for attackers. Bitcoin's security instead only guarantees a low probability of success for an attack (which is still profitable in expectation!). This does not change when delta blocks are included, and in my view casts doubt on the validity of the security analysis that is presented.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

The security analysis relies on several approximations which is methodologically flawed. Security is typically analyzed using bounds that represent worst-case assumptions for the defender. Approximations may be in-exact in ways that may invalidate claims.
The bottom line is this: I am not sufficiently convinced by the security proof in the paper, which I consider to be the main technical contribution. I'd also like more clarification on how this scheme fairs in comparison to other similar ones, including IBLTs and other weak block schemes that do not necessarily chain weak blocks internally.

Additional detailed comments:

p1: Bitcoin's tx capacity limited due to miner's hesitation… There are other limitations, including among other things decrease in security for larger blocks, and increased proportional payments to large miners.

p1: the term impedance is used, but I do not know what this means in a networking context. A precise definition is needed, or more standard terminology. Units are of time per byte: time for what event?

p2: Assumption 4: Why would this assumption necessarily hold? If we ignore the current arbitrary 1MB limit, the protocol enforced block size is a security measure (at whatever size it will eventually settle). Why is free market equilibrium necessarily smaller? Supporting claims / intuitions/ evidence needed here.

Later in the paper (sec 7), the fast block approximation is derived. How does this mesh with assumption 4? If blocks are fast, orphan rates are low, and miners will increase block size (until system limit is hit, or until blocks are no longer propagating fast rel. to block rate).

p2: "…have argued that fees that result from orphaning risk…" This is unclear. How do fees result from orphaning risk? I cannot follow the claim.

Sec 2: list of symbols: Some of these need to be defined more precisely. E.g., what is the "orphaning risk incurred at start of double-spend attack" M_0? Cost of a double spend attack to whom? The impedence (z) appears here again, but I am still left wondering what it is. (The time to propagate blocks to other miners depends on the structure of the network. Decker et. al. collected data on this in various works and have shown that the time at which nodes receive a certain block varies. In fact, some small number of outliers always receive the block much later than most other nodes)

I would much prefer a model section with a clearer set of definitions and assumptions. Some of the terms are explained later in the paper, but not nearly formal enough.

Fig 1: a bit misleading. The number of delta blocks per regular block is not constant, but rather random (Poisson dist.). The circles in the figure are not very informative. What is the chain structure? Which block points to which?

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

Fig 3: There is some mathematical model that underlies this figure. It should be clearly stated and statements proven (assumptions on concavity / convexity of functions, etc) so that we can understand if the claim is reasonable. The drawing alone is not sufficiently clear or precise in explaining the underlying assumptions.

Note 20: is the small miner approximation appropriate? This should be clearly stated in the model section and not hidden in a note. I tried to come up with an explanation to the formula this note refers to (rho supply) and ended up tracing this back to a previous (unpublished paper) by the author, where again it is based on an approximation for the orphaning risk of a node attributed to Andresen. This again leads to a github note by Andresen where no explanation of the formula is given. Please provide a clear derivation of this. What are the underlying assumptions? What is the error term in the approximation?

Section 6: There seems to be something close to a definition of the impedance here (tau = z \Delta Q + tau_0), or at least its relation to propagation time (propagation to 50% of nodes? 100%? Unclear). Is this the definition?

Section 8: This section analyzes the security of zero-confirmation transactions.
Because of the existence of delta blocks, I think the term zero confirmation is no longer obvious and needs a bit more clarification.

The attack that is considered here is only a form of double spending of the weak blocks. The section opens with the statement: "To double spend a transaction… an attacker must produce a weak block with greater fees…". Why is this the only possible attack? I think other approaches also need to be analyzed including a Finney attack with regular blocks. Miners can additionally include transactions of their own with added fees to increase the weight of their delta blocks. A single weak block thus loaded with sufficiently high fees can override a longer chain created by the network, but also a somewhat shorter chain can be augmented with these extra fees. What is the cost of this? Whatever it is, it will be profitable given a sufficiently high transaction that is being double-spent.

The attacker is assumed to have the same orphaning risk as the honest nodes. Is this reasonable? What if he invests more in communication infrastructure?

"For the attacker we cannot use the fast block approximation" Why? Can't the attacker continually send his delta blocks but keep the last part of his chain secret? I think this statement needs further explanation. In particular, the protocol needs to be exactly explained w.r.t how it deals with branches that are off the chain (the bitcoin protocol for example saves off-chain blocks, in case they do eventually turn out to be the longest chain)

Section 9: this is interesting. How does it change the analysis? At some point a deeply nested subchain will no longer uphold the fast block approximation. How deep down is it safe to proceed?

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

1B. Author's Response

Dear Editor:

Firstly, I would like to thank you and the three thoughtful reviewers who reviewed my paper. After addressing the reviewer comments, the revised version of my paper—that I submit to you today—is much stronger.

The common theme among all reviewers was that my explanation for why subchains provide double-spend security for unconfirmed transactions was weak. To address the weaknesses pointed out by the reviewers, I have completely re-written the section on double-spend security (which is now Section 7).

In this new security analysis, I categorize the miners as either default-compliant or petty-compliant [1]. Default compliant miners follow the specified protocol and always mine on top of the longest subchain. Petty compliant miners will deviate from the protocol to facilitate double spend attacks if doing so is profitable. With this new framework, I think the new explanation I give in Section 7 makes it clear how subchains add security to unconfirmed transactions.

I address each of the reviewers' comments point-by point below.

[1] Carlsten, M., Kalodner, H., Weinberg, S. M., Narayanan, A. "On the Instability of Bitcoin Without the Block Reward." ACM CCS 2016.
http://randomwalker.info/publications/mining_CCS.pdf

Reviewer A:

Thank you for this paper.

I've not been able to comment on the math, due to my lack of math-skills. But I do understand the reasoning and conclusions. I hope that's fine.

The main purpose of the paper is to reduce orphaning risks and to make 0-conf double spends less probable to succeed. The Orphaning risk is reduced by minimizing the amount of data sent at block propagation time. The 0-conf security is increased by miners' incentives to stay on the subchain with the most fees.

The idea of weak blocks is very interesting. The obvious benefit from subchains (and other weak block proposals) is reduced orphaning risks. But the less obvious (that I hadn't though of before) effect is that 0-conf transactions may be safer, but I think the paper needs to be more convincing on that matter since the paper lacks elaboration on different types of double spend attacks.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

Major comments:

In section 1, one of the assumptions made is that "Miners are rational, short-term profit-maximizing agents". I'm missing some elaboration on how probable this is to be true, both today and in the future. Today, many miners use default settings of bitcoind, meaning they do what the developers considers best for the network as a whole, but when margins tightens, miners might be forced into tweaking their policies and possibly make agreements with each other and with certain payment processors, merchants, exchanges and so on.

- This is a pretty standard assumption for economics modeling that's related to the "perfect competition" concept. Although the assumption is imperfect, it makes the problem mathematically tractable. I changed the wording here to say that I'm assuming "perfect competition" and cite the Wikipedia entry on the concept.

- Furthermore, I have split the miners into two groups: those that follow the protocol obediently (default compliant), and those that will disobey the protocol to facilitate double-spend attacks if doing so is profitable (petty compliant).

I think the incentives for miners to propagate its weak blocks are unclear and needs more elaboration. For example, would it really be more profitable for me to publish a weak block if I find one? If I find the next strong block, I would benefit from it only if it's immediately after my weak block. If there's another weak block after my block, it made no difference at all. On the other hand, if I don't find the next strong block, I've just given away fast block propagation, ie low orphan risk, and the fee benefit of weak blocks, to someone else.

- As you said, you benefit from publishing a weak block if you find the next strong block. But since *you can't know ahead of time* whether or not you will find the next strong block, you are better off in expectation to publish all weak blocks that you find. I have improved Section 3 – 5 to make the benefit to miners more clear and added the sentence at the end of Section 5: *"Miners are naturally incentivized to share each Δ-block they find, as doing so reduces the orphaning risk of their candidate block."*

Section 8: This section assumes that an attacker will have to (the word "must" in the first sentence of the section) increase its weak block size, by adding transactions (for the fees), faster than the rest of the network in order to produce a higher-fee subchain. But the main goal for the attacker is to perform a successful double spend. Assuming all miners (but the attacker) will build off the highest-fee subchain, all the attacker have to do is create a weak block with ONLY the double spend, but with a higher fee than the highest-fee subchain. If he succeeds, all the hashing power will switch over to his double spend subchain. If he fails, and the next weak block is built off of the honest chain, the attacker will simply create a new block template with a new double spend with the fee increased to be bigger than the new highest-fee subchain.

In a sense, double spending becomes easier, because the attacker have several chances to succeed: one per weak block produced in the honest subchain. To minimize the fee of the

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

double spend, the attacker can of course add as many other transactions he wishes as long as they don't make the block too big (with respect to orphaning risk) or conflict with his subchain.

- As explained in my cover letter, I have re-written the section on the security of unconfirmed transactions. I believe it now answers all of these concerns.

Section 8: If my previous comment is valid, then much of section 8 (the statement that it helps 0-confiramtion transaction security) is invalid for high value double spends. For low value double spends it might still work as described in the paper, because there is no sense in setting the fee higher than the value of the payment. The calculus must take into consideration all reasonable types of double spend attacks.

- As explained in my cover letter, I have re-written the section on the security of unconfirmed transactions. I believe it now answers all of these concerns.

The assumption (1) in section 1 implies that miners accept replace-by-fee transactions. This means that a transaction can be cheaply replaced (just as today) until it's included in a weak block of the highest-fee subchain. When it's included in the highest-fee subchain the cost of replacing it depends on how deep (fee-wise) in the subchain it is buried, and if that cost increase is bigger than the fee increase of the replacement transaction, it's not worth it.

- Yes.

It seems like RBF is incompatible with subchains since you can only add transactions to a subchain, not remove (or swap) them. In order to replace an opt-in replacable transaction, you must either build a higher-fee subchain or not include replacable transactions at all in the subchain, meaning they are put only in strong blocks. You cannot foresee if your block template is going to be a strong or weak block, so you cannot ever include a replacable transaction without making it unreplacable.

- I have added a paragraph in the conclusion that deals with the side effect on RBF:

- *"Subchains also produce a side effect on the replace-by-fee (RBF) logic incorporated into some Bitcoin clients (e.g., Bitcoin Core). RBF is essentially a tool to make it easier for users to "bribe" miners to swap the first-seen version of a transaction with a double-spent version. However, rather than facilitating fraud, RBF's stated aim is to provide a means for users to "unstick" transactions stuck due to too low a fee. RBF will work unchanged with the proposed subchain technique for transactions that have not yet been included in the longest subchain; however, RBF will no longer work (or will require a much greater "bribe") for transactions that have been included. This is not a problem, however, because in this latter case, the transaction is very likely to be included in the next block anyways, so the user has little reason to bump the transaction's fee."*

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

Minor comments:

- The "List of Symbols" section is very helpful.
- Section 4, first paragraph: "A miner is thus financially incetivized to build off the highest-fee subchain": Maybe not true, a miner may have other transactions it wants to mine that actually conflicts with the highest-fee subchain. It might be that a miner have an agreement with a payment provider or other organization that forces them to select transactions in other ways. The next sentence then falls too "Since all miners have the same incentives...".

    - I've clarified the definition of a miners and made it more clear that we're assuming that some are "default compliant" and others are "petty compliant."

Reviewer B:

This paper deals with the idea of subchains. Subchains are composed of weak blocks. Weak blocks are solved by miners in the same way (strong) blocks are solved, only with a weaker (larger) target. The advantage of using subchains is that miners can progressively send smaller pieces of information to the network when building weak blocks. Hence, this has an important impact on the risk of having a weak block orphaned and hence lowers the cost of solving blocks, allowing more transactions to be included in strong blocks. The author explains the mechanism of subchains very clearly and incentives faced by miners when using it (Section 3-4-5), formally shows their advantage in terms of orphan risk and hence in terms of modifying the block space offer function (Section 6), formally computes the cost of trying to outrace the network for an attacker who would like to double-spend some coins (Section 8) and finally notes the possibility to extend this idea of subchains by nesting them (Section 9).

This paper is very interesting and of great interest for the crypto-currency community. I recommend that Ledger publishes this article. I still have a few comments that I would like the author to address.

Major remarks:

One of the major advantages of subchains that the author states in the introduction and the in conclusion is that it does not require a fork (soft or hard) in order to be implemented. It is true that, as it is presented, it only requires participating miners to agree on this protocol to understand each other. And as noted "it does require participation from a significant fraction of the network hash power in order to be useful." The study is then carried on assuming the whole set of miners agree on the subchain design. However, none of the transition (from a situation in which only a small fraction miners adopt subchains) period is discussed. I am not sure this period is only a period of unusefulness of subchains as suggested in the above quoted sentence. Instead, I suspect that there are negative incentives to join the subchain movement. For instance, if only a small fraction of the miners agree on the subchain design (I am aware it is not the question the author adresses), a) either they have to broadcast at some point their strong blocks to the rest of the network or b) they don't. In case a), the fact that strong blocks built with subchains are larger than "regular" strong blocks becomes a disadvantage for the

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

miners having adopted the subchain design ("regular" strong blocks miners are supposed to have the optimal number of transactions in their block).

- Subchains are only marginally larger due to the fixed-byte size reference to the subchain's tip (perhaps 64 bytes larger). That said, in the "transition phase" miners would probably want to follow both the subchain protocol and the existing protocol, and this extra work might be a disadvantage, as you point out. I have added a mention of this in the conclusion:

- *"…although miners would benefit by using the subchain technique if other miners also used it, during the "bootstrapping" phase before the protocol is widely deployed, supporting both standard block propagation and the subchain technique may impose a net cost on forward-thinking miners. How we would move from the current block propagation regime to the more efficient subchain regime is not clear."*

Of course, this would mean that this effect should be anticipated and the consequence are not clear (and certainly depends on who has an incentive to broadcast the strong block). In case b), a fork actually occurs. It is not a fork in the sense of a software change but it is still a fork with two incompatible blockchains. Then, at best, there is a coordination game between miners. This is a qualitative and certainly not complete reasoning but I would suggest the author addresses this question maybe by discussing it in the conclusion a little bit more extensively and clearify the point of a mere unusefulness of subchains when the participation of only a small fraction of the hash rate is "voting" for subchains (maybe only to say that it is not clear if the transition is feasible).

- See comment above.

In the whole paper, the author states that one of the main advantages of subchains is to increase the number of transactions processed by Bitcoin. (eg p1: "Unlike Visa, Bitcoin's transactional capacity is limited due to miners' hesitation to produce blocks containing large volumes of new transactions.") In the formal study, it is transparently stated as an assumption that "The free-market equilibrium block size is smaller than the protocol-enforced block size limit (if such a limit exists)." However, today, the problem of limited number of transactions is more a problem of protocol-enforced block size limit than a orphaning risk driven fee problem. Maybe the "BLOCKSIZE LIMIT DEBATE WORKING PAPER" header does not help. In my opinion, this paper is more a proposition for a cost-cutting production function (when the blocksize limit is not binding) rather than bringing arguments in the blocksize debate. This should be more clear in this article as I guess today's context is very present in today's readers' minds (even though I understand this blocksize limit might be only a temporary issue for Bitcoin).

- Yes, the running header should have been removed.

- I have added a note to the conclusion that the results and model would be less representative of reality if "blocks are always full" in the future.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

Minor remarks:

p3: Fig1 is quite misleading because it lets the reader think that the number of weak blocks per strong block is fixed whereas it is not the case. I would just remove this figure as the text is explanatory enough in my opinion.

- I've added a note to the figure caption to explain that this is an "idealized representation."

p7: I am not sure about the τ_0 term in $P_{orphan}=1-e^{\frac{\tau_0}{T}}e^{zQ\DeltaT}{T^2}$ equation. Indeed, if all miners need τ_0 to spread their solution to the majority of other miners(assuming that all other miners mine empty blocks) weak or strong, then, the probability to be considered should be the one that no other miner finds a block between 0 and $\tau - \tau_0$. Maybe Figure 5 should be modified accordingly. Otherwise it means that other miners have instant spreading of their block.

- I think I deal with tau_0 correctly. If everyone mines empty blocks, there will still be orphan races due to the tau_0 term. This is reflected in Fig. 5 with "latency limit" region shaded in purple.

p7: it should be made clear that another assumption is used here: All miners assume that other miners mine empty blocks. This has important implications in game theory (see Houy's "The Bitcoin Mining Game").

- I have removed the model from my fee market paper and now just rely on the cost for block space obeying the law of demand. This generalizes the result and makes the analysis simpler.

- I have also added a citation to Houy's "The Bitcoin Mining Game."

Fig 5: The considered size per transaction is not specified in the text. It should be for the reader to understand the link between the upper and lower scales.

- I have added a second scale to shown "transactions per second" in addition to the block size.

p8: "the rate at which orphaning risk, M, is incurred". It is not clear here what the definitions are (for those who did not read the "market fee exists" paper by the same author). One more line defining M would certainly be welcome. (maybe even in the beginning of Section 4).

- As explained in my cover letter, I have re-written the section on the security of unconfirmed transactions. I believe it now answers all of these concerns.

Fig8: same as Figure1: it is too "symmetrical" for not being misleading. Again, the text is explanatory enough in my opinion.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

footnote 24: "Since ?supply = zRT" should be "Since ?supply = zRT^{-1}".

Equation 10: If I am not wrong, as it is proved, Equation 10 is obtained by integrating until infinity an expression that is a power series about 0. Maybe table 1 could just be obtained by means of numerical computation rather than trying to have an explicit approximate expression (as said above, maybe not that approximate as it is).

- As explained in my cover letter, I have re-written the section on the security of unconfirmed transactions. I believe it now answers all of these concerns.

Reviewer C:

The paper is generally nicely written and clearly a great deal of work has gone into presentation, figures, etc. Still, I found several problems with the analysis in the paper.

The analysis assumes all miners have similar mempools, and are working on the same block. This may be true when attackers are not present, but what if the attacker sends many conflicting transactions to different nodes in an attempt to increase variations between mempools? How does the protocol react then?

- I have clarified in Section 3 that transactions that are verified in a subchain (weak block) take precedence over conflicts in a node's mempool:

- *"For conflicting (double-spent) transactions, the transaction verified in a subchain has priority over one only admitted into mempool. Note that this behavior represents a departure from the Satoshi protocol where miners will only replace transactions in mempool if a conflicting transaction is included in a strong block (subchains extend this behavior to weak blocks too). This departure is necessary so that miners, under normal conditions, converge upon a single subchain."*

Why would the first delta block be small, and not the size of an entire block? Miners should usually have something in the mempool that they can start hashing that would be more profitable than a small delta block.

- It's not necessarily small. Rational miners will include all TXs that pay a fee greater than the marginal orphaning risk. I think this is already sufficiently clear from the explanation and diagrams in (what is now) Section 5.

There seems to be no accounting for the arrival of transactions with varying fees. If a high paying transaction suddenly appears in a delta block, miners may want to include it, but also to evict previously included low-fee transactions that had been in the previous version of the block being processed.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

- It's implicitly assumed that the delta block is always changing such that it contains the TXs that maximize the expectation value of the miners profit. I have added the sentence in Section 5: *"Note that miners will dynamically adjust their block candidates as new transactions enter mempool, to maximize expected profits."*

The security analysis relies on establishing that an attacker will suffer a cost from attacking. This is highly problematic, as it is already well established that double-spending attacks in Bitcoin are profitable for attackers.

- This work doesn't suggest otherwise. It only shows that there's a cost associated with double-spending a TX that's been included in a subchain. Attempting the double-spend would still be profitable for if the attack cost is less than the bribe paid.

- What's important is that without subchains, the cost of the attack is zero as miners can freely replace a transaction in mempool with a later double-spend version. With subchains, there is now a cost to doing so because the miner can no longer take advantage of the pre-propagated transaction fees.

- As explained in my cover letter, I have re-written the section on the security of unconfirmed transactions. I believe it now answers all of these concerns.

Any miner engaging in selfish mining occasionally creates long secret chains that can be used to double spend. This strategy is profitable for attackers. Bitcoin's security instead only guarantees a low probability of success for an attack (which is still profitable in expectation!). This does not change when delta blocks are included, and in my view casts doubt on the validity of the security analysis that is presented.

- What happens during edge cases like selfish mining are outside the scope of the paper, as the paper already covers a lot of material.

The security analysis relies on several approximations which is methodologically flawed. Security is typically analyzed using bounds that represent worst-case assumptions for the defender. Approximations may be in-exact in ways that may invalidate claims.

The bottom line is this: I am not sufficiently convinced by the security proof in the paper, which I consider to be the main technical contribution. I'd also like more clarification on how this scheme fairs in comparison to other similar ones, including IBLTs and other weak block schemes that do not necessarily chain weak blocks internally.

- I have added a "Related Work" section (Section 9) that compares subchains to other similar ideas.

Additional detailed comments:

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

p1: Bitcoin's tx capacity limited due to miner's hesitation… There are other limitations, including among other things decrease in security for larger blocks, and increased proportional payments to large miners.

- I have added the words "in part" to that sentence:

- *"Unlike Visa, Bitcoin's transactional capacity is limited in part due to miners' hesitation to produce blocks containing large volumes of new transactions"*

p1: the term impedance is used, but I do not know what this means in a networking context. A precise definition is needed, or more standard terminology. Units are of time per byte: time for what event?

- I have clarified this by explicitly stating the model used in the introduction: *"Information propagates from the miner who solves a block to the other miners according to the simplified model $\tau = \tau\_0+zQ$, where τ is the propagation time, Q is the number of bytes propagated, and z and τ_0 are empirical constants."*

p2: Assumption 4: Why would this assumption necessarily hold? If we ignore the current arbitrary 1MB limit, the protocol enforced block size is a security measure (at whatever size it will eventually settle). Why is free market equilibrium necessarily smaller? Supporting claims / intuitions/ evidence needed here.

- This assumption might not hold in the future, especially if the 1 MB limit is never raised. I have addresses this in the conclusion with the following paragraph:

- *"We also assumed that the protocol-enforced block size limit (if one exists) was greater than the free-market equilibrium block sizes produced. This was the regime that Bitcoin was operating under from January 2009 until mid 2015. If the network continues operating in a saturated-block regime as it is today, the marginal orphaning risk for a given transaction could be significantly less than that transaction's fee, and so the benefit to miners of cooperating to build subchains would be significantly reduced. The double-spend resistance of unconfirmed transactions would likewise be reduced."*

Later in the paper (sec 7), the fast block approximation is derived. How does this mesh with assumption 4? If blocks are fast, orphan rates are low, and miners will increase block size (until system limit is hit, or until blocks are no longer propagating fast rel. to block rate).

- I have removed this model entirely as it was unnecessary. All I assume now is that the block space satisfies the "law of demand"—that is, that supply curve is a monotonically increasing function of the block size. This generalizes the results of the paper and makes it easier to understand as well.

p2: "…have argued that fees that result from orphaning risk…" This is unclear. How do fees result from orphaning risk? I cannot follow the claim.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

- I have added the following paragraph to the beginning of Section 5 to help clarify:

- *"With conventional block propagation, a miner must balance the additional fee revenue he earns by making his block bigger, with the decreased orphan risk he enjoys by making his block smaller. The free-market equilibrium block size is the point where a smaller block would result in a smaller expected profit due to too many fees left in mempool, while a larger block would also result in a smaller expected profit due to too high an orphan risk. The author describes this equilibrium in detail in his paper on Bitcoin's transaction fee market.[ref]"*

Sec 2: list of symbols: Some of these need to be defined more precisely. E.g., what is the "orphaning risk incurred at start of double-spend attack" M_0? Cost of a double spend attack to whom? The impedence (z) appears here again, but I am still left wondering what it is. (The time to propagate blocks to other miners depends on the structure of the network. Decker et. al. collected data on this in various works and have shown that the time at which nodes receive a certain block varies. In fact, some small number of outliers always receive the block much later than most other nodes). I would much prefer a model section with a clearer set of definitions and assumptions. Some of the terms are explained later in the paper, but not nearly formal enough.

- As explained in my cover letter, I have re-written the section on the security of unconfirmed transactions. I believe it now answers all of these concerns.

Fig 1: a bit misleading. The number of delta blocks per regular block is not constant, but rather random (Poisson dist.). The circles in the figure are not very informative. What is the chain structure? Which block points to which?

- I have added an additional sentence to explain that this is an "idealized" diagram and the in reality, there may be more or less weak blocks per strong blocks—the picture just represents "the average."

Fig 3: There is some mathematical model that underlies this figure. It should be clearly stated and statements proven (assumptions on concavity / convexity of functions, etc) so that we can understand if the claim is reasonable. The drawing alone is not sufficiently clear or precise in explaining the underlying assumptions.

- I have clarified why the curves have the concavities they do in Section 5:

- *"The expected cost, ⟨C⟩, associated with this risk is depicted in Fig. 4a as a function of block size. By assuming only that block space obeys the law of supply,[ref] it follows that this curve is superlinear in Q, although the concavity can also be deduced using technical arguments.[ref] The second curve represents the maximum fees, F, available from transactions in mempool for a block of size Q. It follows, by definition, that this curve is sublinear.[ref]*

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

Note 20: is the small miner approximation appropriate? This should be clearly stated in the model section and not hidden in a note. I tried to come up with an explanation to the formula this note refers to (rho supply) and ended up tracing this back to a previous (unpublished paper) by the author, where again it is based on an approximation for the orphaning risk of a node attributed to Andresen. This again leads to a github note by Andresen where no explanation of the formula is given. Please provide a clear derivation of this. What are the underlying assumptions? What is the error term in the approximation?

- I have removed this entirely from the analysis, as this level of detail was not actually needed.

Section 6: There seems to be something close to a definition of the impedance here (tau = z \Delta Q + tau_0), or at least its relation to propagation time (propagation to 50% of nodes? 100%? Unclear). Is this the definition?

- I have clarified this by explicitly stating the model used in the introduction: *"Information propagates from the miner who solves a block to the other miners according to the simplified model $\tau = \tau_0+zQ$, where τ is the propagation time, Q is the number of bytes propagated, and z and τ_0 are empirical constants."*

Section 8: This section analyzes the security of zero-confirmation transactions. Because of the existence of delta blocks, I think the term zero confirmation is no longer obvious and needs a bit more clarification.

- I now ensure I use the word "confirmation" for strong blocks, and "verifications" for weak blocks. I also use "unconfirmed TX" rather than "zero-confirm TX" whenever possible.

The attack that is considered here is only a form of double spending of the weak blocks. The section opens with the statement: "To double spend a transaction… an attacker must produce a weak block with greater fees…". Why is this the only possible attack? I think other approaches also need to be analyzed including a Finney attack with regular blocks. Miners can additionally include transactions of their own with added fees to increase the weight of their delta blocks.

A single weak block thus loaded with sufficiently high fees can override a longer chain created by the network, but also a somewhat shorter chain can be augmented with these extra fees. What is the cost of this? Whatever it is, it will be profitable given a sufficiently high transaction that is being double-spent.

The attacker is assumed to have the same orphaning risk as the honest nodes. Is this reasonable? What if he invests more in communication infrastructure?

"For the attacker we cannot use the fast block approximation" Why? Can't the attacker continually send his delta blocks but keep the last part of his chain secret? I think this statement needs further explanation. In particular, the protocol needs to be exactly explained

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2016.40

w.r.t how it deals with branches that are off the chain (the bitcoin protocol for example saves off-chain blocks, in case they do eventually turn out to be the longest chain)

- These are all valid comments. As explained in my cover letter, I have re-written the section on the security of unconfirmed transactions. I believe it now answers all of these concerns.

Section 9: this is interesting. How does it change the analysis? At some point a deeply nested subchain will no longer uphold the fast block approximation. How deep down is it safe to proceed?

- Subchaining is possible until propagation time is limited by the y-intercept of the "propagation time vs size" curve and not so much by the slope of the curve. I think this is sufficiently clear here and in the section on orphaning risk.

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.2016.40