LEDGER
ledgerjournal.org

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.40

RESEARCH ARTICLE

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

Peter R. Rizun*†

Abstract. Orphan risk for large blocks limits Bitcoin's transactional capacity while the lack of secure instant transactions restricts its usability. Progress on either front would help spur adoption. This paper considers a technique for using fractional-difficulty blocks (weak blocks) to build subchains bridging adjacent pairs of real blocks. Subchains reduce orphan risk by propagating blocks layer-by-layer over the entire block interval, rather than all at once when the proof-of-work is solved. Each new layer of transactions helps to secure the transactions included in lower layers, even though none of the transactions have been confirmed in a real block. Miners are incentivized to cooperate building subchains in order to process more transactions per second (thereby claiming more fee revenue) without incurring additional orphan risk. The use of subchains also diverts fee revenue towards network hash power rather than dripping it out of the system to pay for orphaned blocks. By nesting subchains, weak block verification times approaching the theoretical limits imposed by speed-of-light constraints would become possible with future technology improvements. As subchains are built on top of the existing Bitcoin protocol, their implementation does not require any changes to Bitcoin's consensus rules.

1. Introduction

Bitcoin's performance as a payment network is hardly impressive. In 2015, it processed an average of 1.4 transactions per second while merchants waited on average eight minutes to receive initial verification from a miner that a transaction would likely be included in the permanent Blockchain ledger.[1, 2] In contrast, the Visa network processed over 2,000 transactions per second,[3] and—with chip-and-PIN technology—merchants received authorization and PIN-verification in under a second.[4] Unlike Visa, Bitcoin's transactional capacity is limited in part due to miners' hesitation to produce blocks containing large volumes of new transactions.[5] Such blocks propagate across the network slowly,[6, 11] increasing the chances that the block is orphaned and the miner's reward is lost.[7] Also unlike Visa, the initial verification of a transaction by a miner is delayed because blocks are propagated on average only every ten minutes,[8] rather than at a rate dynamically tuned to the bandwidth and latency of the network. In this paper, we present a scaling technique called *subchains* to build blocks layer-by-layer—at a small fraction of Bitcoin's ten-minute block time—thereby reducing both orphaning risk and the wait-time for the first verification of a transaction by a miner.

Throughout this paper, we make certain simplifying assumptions. In particular, we assume that:

(Asm. 1) 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.[6, 9, 10, 11] (Block validation time is assumed to be included in the propagation time.)

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

(Asm. 2) The market for block space is one of *perfect competition*.[12]

(Asm. 3) The protocol-enforced block size limit—if such a limit exists—is greater than the free-market equilibrium block size. That is, the block size is constrained by natural supply and demand, rather than by a production quota.

(Asm. 4) The network consists of *default-compliant* miners who reliably follow the agreed-upon protocol and *petty-compliant* miners who will deviate from the protocol to facilitate double-spend attacks if such behavior is profitable.[13] The total hash power controlled by petty-compliant miners is $\chi \ll 50\%$.

This paper makes the following contributions:

Contribution 1: Description of the subchain technique. In Section 3, we describe the subchain technique,[14] which is a practical application of weak blocks[15, 16, 17, 18] that provides incentives for miners to cooperate for the mutual benefit of the network. Its implementation requires neither a hard nor soft fork—but it does require participation from a significant fraction of the network hash power in order to be useful. In Section 9, we illustrate how subchains can be nested, creating a fractal-like blockchain structure where transactions are processed almost continuously.

Contribution 2: Reduced orphan risk. A significant advantage of the subchain technique is revealed in Section 4, when we show how the technique considerably reduces orphan risk for a given sized block. The reduced orphan risk is due to the fact that the block is built layer-by-layer over the ten-minute block interval, rather than propagated all at once the moment the proof-of-work is solved. Using subchains, miners can cooperate to process more transactions per second for a given level of orphan risk.

Contribution 3: Existence of a fee market. We move onto the economics of the transaction fee market in Section 5. We show that although a miner can include all of the subchain's transactions in his block candidate—and thus all of the subchain's fees—without incurring additional orphan risk, he still incurs extra orphan risk for *new* transactions included in his block candidate. This property drives a fee market and economically restricts the rate of subchain growth.

Contribution 4: Proof-of-work security from fee revenue. Certain investigators have argued that fees that result from orphan risk cannot contribute to network security.[19] With a simple diagram, we prove this line of reasoning false in Section 6 by showing that the fees already included in the subchain contribute *directly* to network security in the same way that the block reward does. Only the fees in the new transactions added on top of the subchain go to cover the orphan risk for those transactions. The *total fees* in a typical block are thus much larger than the block's *total orphan risk*.

Contribution 5: Security for unconfirmed transactions. In Section 7 we review the double-spend security offered by unconfirmed transactions under standard block propagation rules, and explain why a lower bound on that security is zero. We then show that double-spending a transaction verified in a subchain has an objectively-measurable cost commensurate with the total fees that have accumulated in the subchain above the transaction an attacker is attempting to double-spend.

2. List of Symbols

For the remainder of this manuscript, the following symbols have the specified meanings.

$\langle C \rangle$	expected cost due to orphan risk	t	time
F	fees	X	subchain factor (weak blocks

			expected per strong block)
P_{orphan}	probability of an orphan race	z	propagation impedance
Q	block size or block space in bytes	χ	hash power controlled by petty-compliant miners
ΔQ	size of Δ-block	τ	propagation time
T	block interval (10 min target)	τ_0	network latency
ΔT	Δ-block (weak block) interval	$\Delta\tau$	propagation time minus latency

3. Weak Blocks and Subchains

To append a new block to the Blockchain, a miner must find a valid proof-of-work. This entails finding a nonce that when hashed together with the previous block's hash and the root hash for the block's transactions, results in an integer less than the network difficulty target.[20] We define a *weak block* as a block that satisfies the weaker requirement

$$\text{hash(previous hash, nonce, root hash)} < \text{weak target},$$

where the weak target is larger than the difficulty target. More plainly, a weak block is a block with enough proof-of-work to be hard to find, but not enough to be a real block. By sharing these weak blocks, miners can cooperate to build *subchains* (Fig. 1).

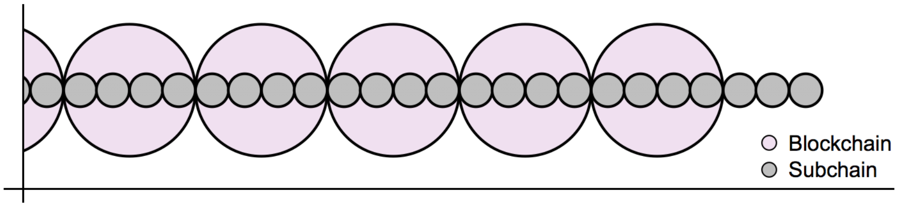

Fig. 1. Miners cooperate to build subchains in order to process more transactions and claim more fees without incurring additional orphaning risk. This illustration visualizes ¼-difficulty "idealized" subchains; due to the luck associated with finding a valid proof-of-work, in reality some strong blocks would contain more than four weak blocks and some would contain less.

Upon accepting a (strong) block, miners begin working on creating the next block in the chain by using the hash of the accepted block as the previous hash (Fig. 2a). When a miner finds a proof-of-work that satisfies the weak target, he broadcasts the weak block to the network. After verifying the weak block, each miner modifies the coinbase reward, appends additional transactions to the block if desired, computes the new root hash, and then continues scanning for a valid nonce (Fig. 2b). We will refer to the new information as the miner's *Δ-block* (Fig. 2f). If again a miner finds a proof-of-work that satisfies the weak target, he broadcasts the new weak block by sending only his Δ-block and the hash of the previous weak block. In this manner, miners can cooperate to build the subchain by transmitting only the new information and a hash that references the subchain's tip.

When a miner finds a proof-of-work that meets the strong target (Fig. 2d), he broadcasts it in the same manner he would for a weak block (*i.e.*, by sending only his Δ-block and the hash of the previous weak block). Nodes recognize this as a valid (strong) block, retain the nonce and coinbase transaction, and close the subchain. The process of constructing a subchain on top of this latest block begins anew (Fig. 2e).

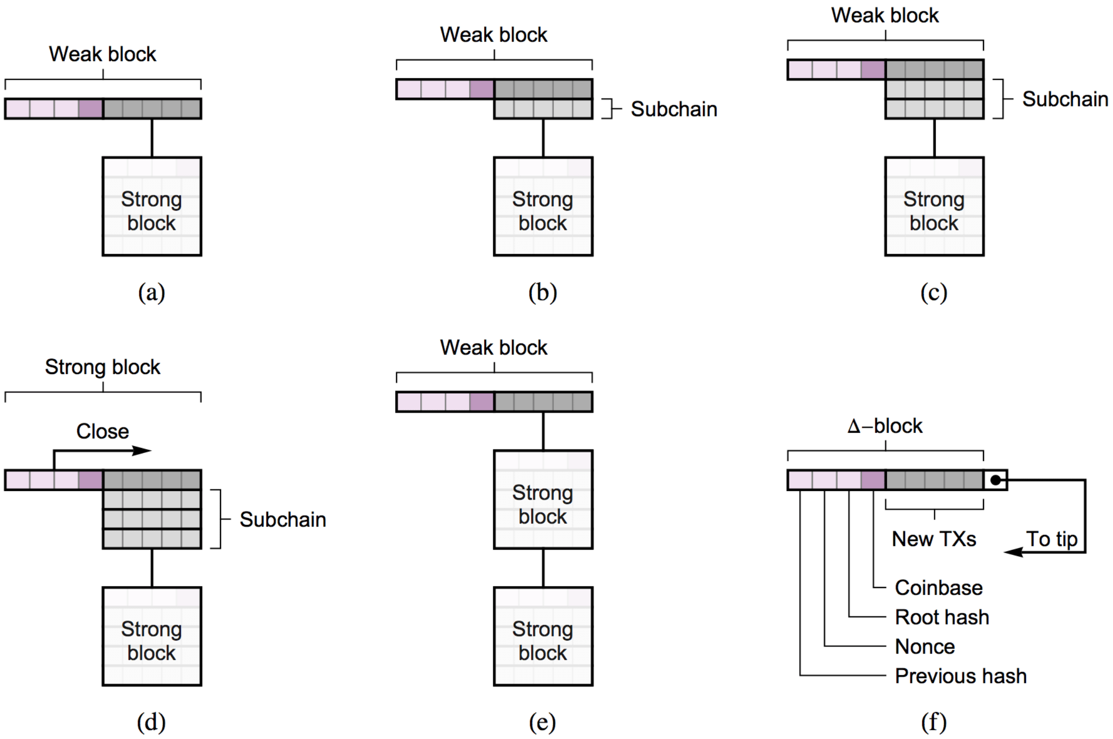

Fig. 2. Miners build subchains layer by layer (a – c), where each new layer corresponds to the solution of a weak block. In (b – d), the subchain (light gray) is common across the network, but the Δ-block is unique to the particular miner. To propagate blocks (weak or strong), miners need only send their Δ-block and a hash that references to the subchain's tip (f), thereby reducing the number of bytes transmitted the moment the proof-of-work is solved. When a nonce that satisfies the strong target is found, the subchain is closed thereby becoming a strong block (d), and miners begin working on a new subchain (e).

If more than a single subchain exists, compliant miners build off the *longest* subchain. In cases where two subchains of equal length exist, miners work on the one they knew about first, switching to the other if it becomes longer. 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. For the remainder of this paper, *mempool* is redefined as the set of transactions that have been neither confirmed in a strong block nor verified in a weak block.

4. Reduced Orphan Risk

A block is orphaned whenever two miners find competing solutions for the next block. The probability of this event depends on how quickly news of a new block spreads across the network. If a block takes time τ to propagate, the probability that the network finds another block during the propagation interval $0 < t < \tau$ is given by

$$P_{\text{orphan}} = \int_0^{\tau} \frac{1}{T} e^{-\frac{t}{T}} dt = 1 - e^{-\frac{\tau}{T}}, \tag{1}$$

where $\frac{1}{T} e^{-\frac{t}{T}}$ is of course the probability distribution for the arrival time of a valid proof-of-work.[21, 22, 23] Subchains reduce a block's propagation time (τ) because only the most recently-added transactions need to be propagated, thereby reducing the probability of orphaning.

Assuming that miners produce equal-sized Δ-blocks with an average period of ΔT, each Δ-block is scaled down by the *subchain factor*, $\frac{T}{\Delta T}$, such that $\Delta Q = \frac{\Delta T}{T} Q$. The propagation time (*cf.* Asm. 1) is thus $\tau = z\Delta Q + \tau_0$, from which it follows that

$$P_{\text{orphan}} = 1 - e^{-\frac{\tau_0}{T}} e^{-\frac{zQ\Delta T}{T^2}}.$$

This equation is plotted in Fig. 3 for various subchain factors and using recent estimates for the network propagation constants ($z = 17$ s/MB and $\tau_0 = 10$ s).[6, 9, 11, 24]

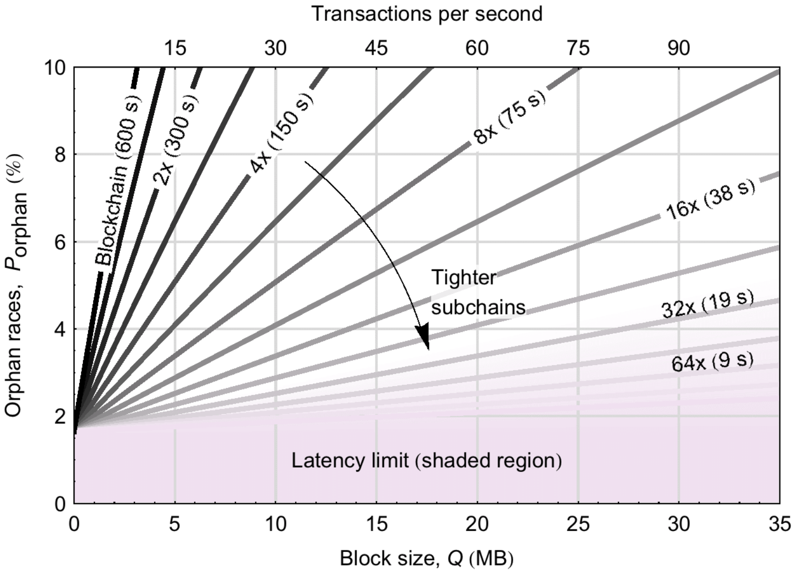

Fig. 3. Subchains help scale Bitcoin by reducing orphan risk for larger block sizes. This chart is based on recent estimates for the network propagation constants ($z = 17$ s/MB and $\tau_0 = 10$ s).[6, 9, 11] For example, a subchain with a target weak block time of 38 s would permit approximately sixteen times more transactions per second at the same level of orphaning risk as without the subchain. The minimum subchain verification time is limited, however, due to network latency (shaded region).

A subchain with $\frac{T}{\Delta T} = X$ would permit approximately X times more transactions per second at the same level of orphaning risk as without the subchain. The minimum useful subchain verification time is limited, however, because the network cannot come to consensus regarding the subchain faster than the network's latency (which, regardless of technology advancements, is limited by the product of the network diameter and the speed of light to approximately 0.1 s).[25]

5. Existence of a Fee Market

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.[26] 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.[5]

This equilibrium changes in one important way in a scenario where subchains are the default mechanism to build and propagate blocks: a miner can now include all of the subchain's transactions in his block candidate—and thus all of its fees—without affecting the block's orphan risk. The reason this is possible is because the miner can now reference the entirety of the subchain with a single hash; the propagation time for that hash does not depend on the size of the subchain that the hash references. The fees in each propagated Δ-block thus add to the subchain's "pot," increasing the effective block reward (as indicated by the black points in Fig. 4a at Q_1, Q_2 and Q_3) but without increasing the expected cost of mining the block.

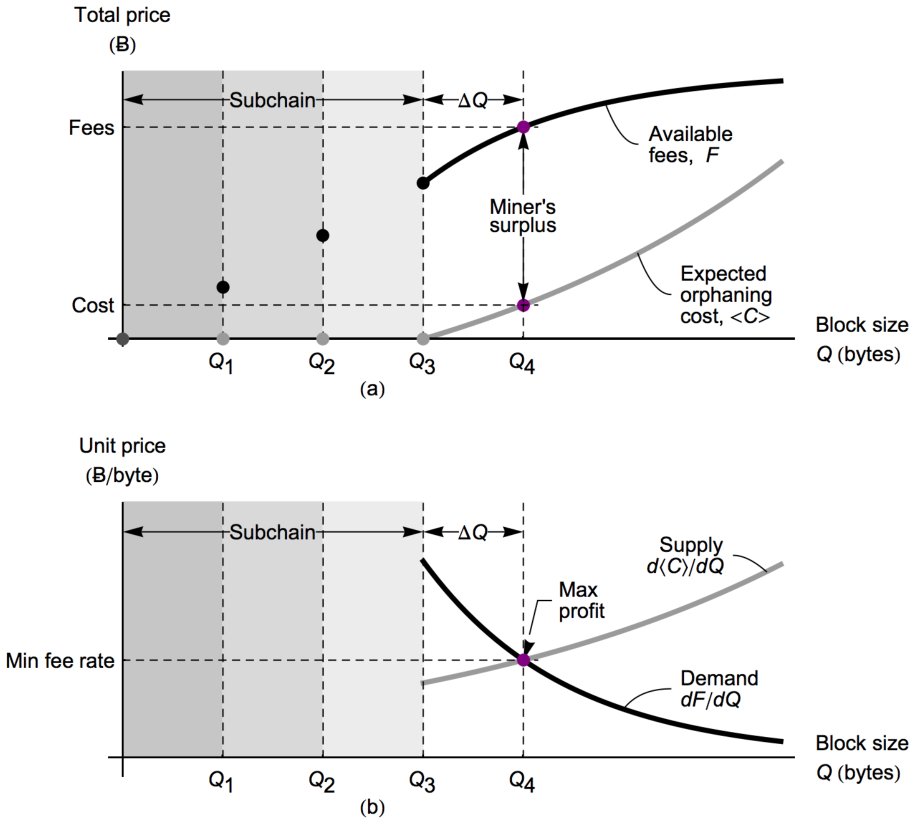

Fig. 4. (a and b) The total fees in a subchain increase each time a Δ-block gets added, effectively "growing the pot." Since a miner incurs no orphaning risk by including the contents of the subchain in his block candidate, he is incentivized to build his Δ-block on top of the highest-fee subchain. However, a miner incurs normal orphaning risk for any *new* transactions he chooses to add. Because of this, he will only include new transactions that pay more in fees per byte than the marginal cost of producing the extra block space to hold those transactions. The block of maximum profit is thus the point where the marginal cost curve ($d\langle C\rangle/dQ$) intersects the marginal fee curve (dF/dQ).

A miner does, however, still incur orphan risk for the *new* transactions included in his Δ-block. The larger he makes his Δ-block, the slower it would propagate across the network, and—in the case where he finds a strong block—the greater the risk he incurs of having his block orphaned and losing the block reward. The expected cost, $\langle C\rangle$, associated with the risk of producing new block space is depicted in Fig. 4a as a function of block size. By assuming only that block space obeys the law of supply,[27] it follows that this curve is superlinear in Q.[28] (The concavity can also be deduced using technical arguments[29]: let c be the value of a valid proof of work and let $P_{\text{orphan}}(Q)$ represent the probability that a block transmitted using Q bytes is orphaned. By including ΔQ bytes of *new* transactions in his Δ-block, a miner is *worse off* by an amount $\langle C\rangle = \frac{c}{1-P_{\text{orphan}}(\Delta Q)} - \frac{c}{1-P_{\text{orphan}}(0)}$ compared to including no new transactions. Using Eq. (1) for the orphaning probability and using the substitution $(\tau-\tau_0) \rightarrow z\Delta Q$ from Assum. 1 gives $\langle C\rangle = ce^{\frac{\tau_0}{T}}\left(e^{\frac{z\Delta Q}{T}}-1\right)$—a superlinear function of Q as expected.) The other curve in Fig. 4a 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.[30] The miner's expected profit is greatest at the block size that maximizes the difference between these two curves, which from elementary calculus occurs at the point where the two curves have equal slopes (*i.e.*, when $d\langle C\rangle/dQ = dF/dQ$), or, translated into the language of economics, at the point where the marginal expected cost is equal to the marginal fee revenue (*cf.* Fig. 4b).[5, 31, 32] Each miner will dynamically adjust his block candidate as new transactions enter mempool, to continually maximize his expected profit. The existence of this equilibrium indicates that a transaction fee market based on orphaning risk exists.

Miners are naturally incentivized to share each Δ-block they find, as doing so reduces the orphan risk for their candidate block.

6. Proof-of-Work Security From Fee Revenue

It is simple to show that fees contribute to proof-of-work security if the network uses the subchain technique (even in the absence of a block size limit). Fig. 5 is a modification of Fig. 4a that considers all of the miner's revenues and costs, including the block reward and electricity for hashing. In a competitive market, the profits for marginal miners will trend to zero. To reconcile this fact with Fig. 5, the total production costs for block space must increase such that the two points marked in purple move closer together. That is, if industry profits were large, miners would tend to deploy more hash power to compete for this profit, thereby shifting the entire production cost curve upwards, increasing hashing costs (due to a rise in network difficulty) and decreasing profits. As shown in Fig. 5, the fee revenue is significantly greater than the orphan risk; this fee revenue—captured Δ-block-by-Δ-block in the subchain—acts no differently than an increase in the block reward would: it serves to increase the network hash rate. We have now shown that fees contribute to proof-of-work security when using the subchain technique.

One subtlety to note is that a miner with revenues and costs as depicted in Fig. 5 would not start to mine until the subchain contained sufficient fees to make the expectation value of his profit positive.[33] Presently, fees are such a small fraction of the block reward that most miners are profitable regardless of fees. However, when total fees are no longer small compared to the block reward, we would expect the instantaneous hash rate to increase every time a new Δ-block (and its fees) is added, as miners with marginally higher electricity costs turn on their machines. Further discussion of this phenomenon is beyond the scope of this paper.

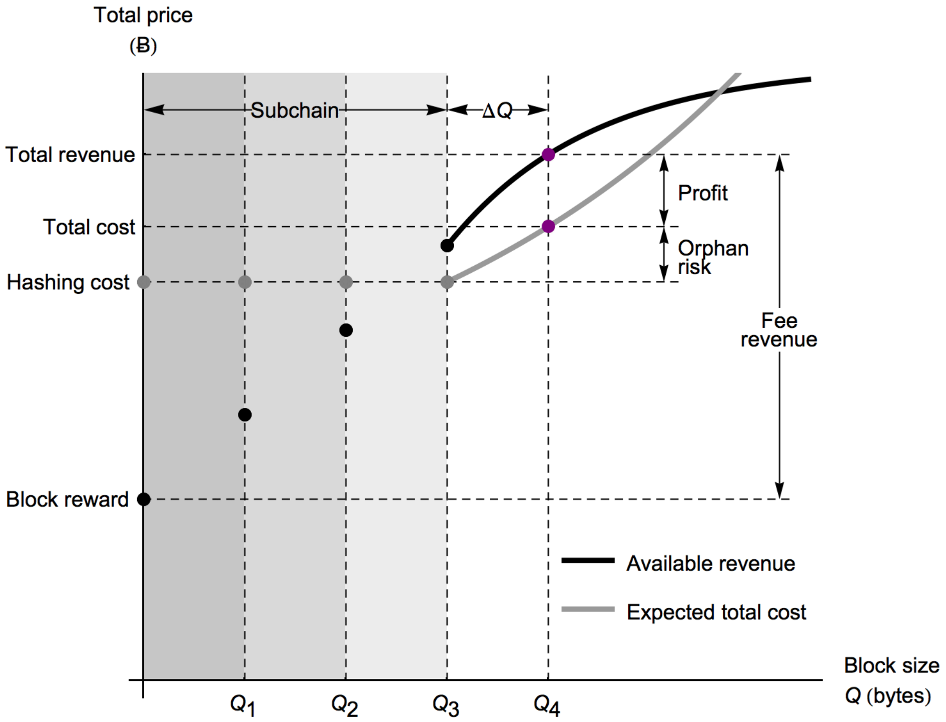

Fig. 5. If excess mining profits are available due to high fees, miners will deploy more hash power to compete for that profit. This causes network difficulty to increase, shifting the curve marked "expected total cost" upwards and reducing industry profits. Fees thus contribute directly to proof of work security by making it more difficult to mine a block.

7. Improved Security for Unconfirmed Transactions

Double-spending today—Consider a scenario where a *scammer* (a dishonest customer) purchases a cup of coffee from a merchant, pays with a bitcoin transaction, and then later tries to double-spend that transaction to reverse his payment. After broadcasting his transaction (and assuming he pays a sufficient fee), nodes relay it across the network, miners incorporate it into their block candidates, it registers with the merchant's listening node, and the merchant hands the cup of coffee to the scammer. Shortly after leaving the merchant's shop with coffee in hand, the scammer then broadcasts the double-spend transaction shown in red in Fig. 6. He attaches a bribe—in the form of a higher transaction fee—trying to entice miners to replace the legitimate transaction with the red transaction, so that the payment to the merchant is never confirmed in a block.

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.40

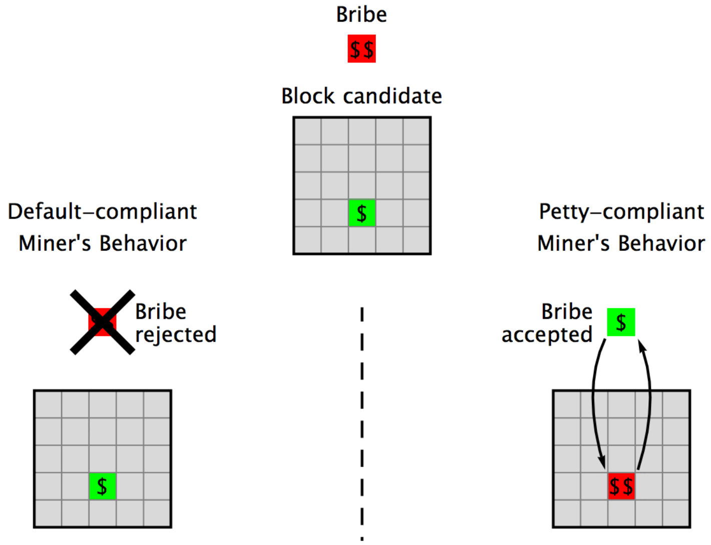

Fig. 6. When offered a bribe to help facilitate a double-spend attack on an unconfirmed transaction, miners can behave in two different ways. Default-compliant miners accept only the first-seen version of a transaction and thus reject the bribe, as specified by the protocol. Petty-compliant miners—recognizing the profit potential of swapping the two transactions—disobey the protocol and accept the bribe. The double-spend succeeds with a probability equal to the fraction of the hash power controlled by petty-compliant miners.

Because the Bitcoin protocol specifies that a miner must accept only the *first-seen* version of a transaction into his mempool, a default-compliant miner would not accept the bribe. On the other hand, a petty-compliant miner—recognizing the profit potential of swapping the two transactions—would. Assuming that petty-compliant miners control a fraction χ of the hash power, this double-spend attack thus succeeds with probability χ (*i.e.*, the probability that a petty-compliant miner finds the next block). Since this petty-compliant behavior remains profitable for even infinitesimally small bribes, it is sometimes said that the lower bound on the security of unconfirmed transactions in Bitcoin is zero.

Double-spending with subchains—To illustrate how subchains improve the security of unconfirmed transactions, we imagine the same double-spend attempt unfolding in a future where miners build blocks using the subchain technique. In this new scenario, when the scammer has left the merchant's shop and broadcasts the double-spend transaction, assume the original green transaction is already verified in a weak block, with additional weak blocks stacked above it, as shown in Fig. 7. Petty-compliant miners can no longer simply swap the red transaction with the green transaction because doing so would break the subchain. To accept the bribe, each petty-compliant miner must now build his block candidate from a layer in the subchain prior to the inclusion of the legitimate (green) transaction. This requirement imposes a real cost on the petty-compliant miner. By following the default strategy, the miner could have included all of the fees in the subchain for zero cost; but by building off an earlier layer in the subchain instead, he either forfeits the fees in the transactions verified in the higher layers or he accepts the added orphaning risk of re-propagating those transactions in the event that he solves the next strong block. To entice this petty-compliant behavior, it follows that the scammer must offer a significant bribe, as explained next.

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.40

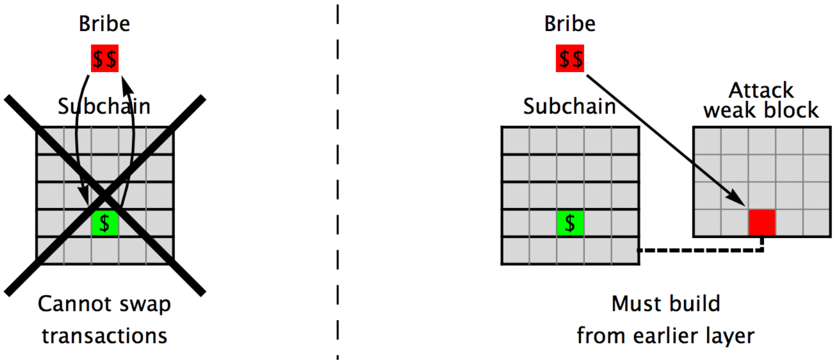

Fig. 7. If default-compliant miners build blocks using the subchain protocol, a petty-compliant miner can no longer swap the green transaction for the red double-spent version without breaking the subchain. To accept the bribe, he must build a block candidate referencing an earlier layer in the subchain. This adds a measurable cost to the double-spend attack.

In a perfectly competitive market, miners include in their block candidates any transaction that pays a fee per byte greater than the expected cost of including the transaction (*i.e.*, greater than the transaction's marginal orphaning risk). Although we will not be rigorous here, it follows that the fee on a *typical* transaction in the subchain will thus be only slightly greater than the expected cost of including that transaction. The cost that the miner bears by mining from an earlier layer in the subchain thus depends on the total fees in the higher layers that he leaves behind, and so the scammer must offer a bribe commensurate with the total fees. The fees in each new weak block found further increase the transaction's security. Despite the increased cost of attempting this double-spend attack, its probability of success remains unchanged and equal to the fraction of the hash power that engages in petty-compliant behavior.

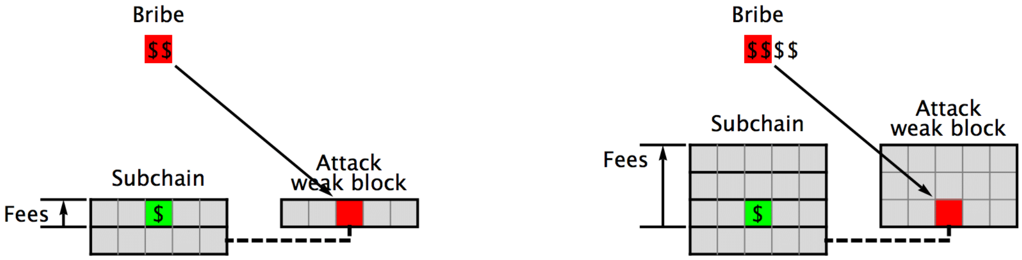

Fig. 8. The cost of the attack depends on the fees above the transaction to be double-spent. As the subchain grows, the miner attempting the attack misses out on being able to claim more and more fees in the pre-propagated transactions, and thus he demands a larger and larger bribe in order to continue the attack.

8. Nested Subchains

Miners must *agree* on the target number of weak blocks per strong block in order for the subchain technique to work. There appears to be a tradeoff in choosing that target: targeting too few weak blocks would lead to higher orphan risk and slower subchain verifications while targeting too many may result in weak blocks found so quickly that convergence upon a single subchain fails. A possible solution to avoid this tradeoff is to use nested subchains (Fig. 9).

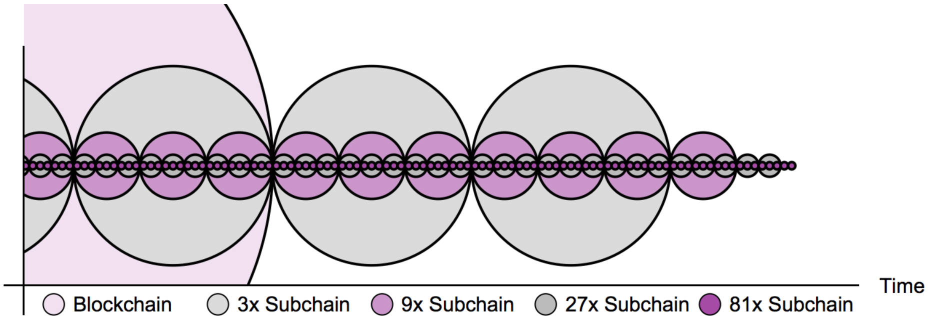

Fig. 9. In a regular subchain, targeting too few weak blocks per strong block would lead to higher orphan risk and slower subchain verifications while targeting too many may result in weak blocks found so quickly that convergence upon a single subchain fails. Subchains can be nested to avoid this tradeoff. This image shows "idealized" trinary nesting of Δ-blocks.

A nested subchain is simply a subchain within a subchain. When a miner finds a block that satisfies the subchain difficulty, the deeper subchain is closed and a new subchain at the deeper level begins. Miners build from the longest subchain at a given nesting depth but ignore longer subchains at deeper nestings. As miner connectivity improves, the subchain verification times can be reduced by adding deeper levels of nesting, fundamentally limited only by the time is takes light to travel across the network.

9. Related Work

Bitcoin's 10-minute block time has long been a frustration for users, spurring developers to create new cryptocurrencies with faster block times. For example, Litecoin has a block time of 2.5 minute while Dogecoin has a block time of 1 minute. Although subchains *decrease* network orphan rates, reducing the inter-block time has the opposite effect of *increasing* them. For example, while Bitcoin's orphan rate is approximately 1%, Litecoin's and Dogecoin's orphan rates (as measured by ProHashing) are 6% and 11%, respectively.[34] Although a finite orphan rate is helpful in driving a fee market, a very large orphan rate has the negative effect of disproportionately benefiting mining pools that control a large fraction of the network hash power.[35] This limits the extent to which faster block times could be used to speed up transaction verification (even if such a change were politically feasible).

In the wake of Bitcoin hitting its "1 MB block size limit" in the summer of 2015, engineers and researchers began working on the problem of scalability in earnest. A simple but highly effective improvement to block propagation was implemented by Peter Tschipper in the Bitcoin Unlimited client. Known as Xtreme Thin Blocks,[36] this Bloom-filter based technique allowed blocks to be transmitted with 24 times fewer bytes and over 5.6 times faster than using conventional block propagation.[37] It can be used together with subchains for the propagation Δ-blocks, providing better scalability than either technique applied on its own.

More radical scalability proposals that would involve changes to Bitcoin's consensus rules have also received substantial interest. Bitcoin-NG separates Bitcoin's conventional blocks into "leader" blocks containing the proof-of-work followed by micro-blocks containing the transactions.[38] Because the transactions are sent after the proof-of-work, orphan risk is not only reduced, it is eliminated entirely.

Another related scalability proposal is known as "braids."[39] Rather than referencing only a single block as a parent (as required by the Bitcoin protocol and orphaning any competing parents), the braid technique allows a miner to reference any number of parents, also potentially eliminating orphans.

It is not however clear that eliminating orphan risk is a net positive. As described in Section 5, it is the small amount of remaining orphan risk that would help drive a transaction fee market if miners were to adopt the subchain protocol. What the transaction fee market would look like if Bitcoin-NG or braids were implemented is not clear.

10. Conclusion

We presented a technique called subchains designed to improve the security of unconfirmed transactions and increase the number of transactions per second the Bitcoin network can process. Subchains are formed as a series of weak blocks, with the next weak block building a new layer of transactions upon the previous weak block. Miners transmit blocks (both weak and strong) by sending only the latest layer of transactions along with a hash that references the previous layers.

The bulk of this paper was dedicated to exploring four important properties that would emerge if the subchain technique were widely deployed by the network's hash power. These properties were:

(Prop. 1) Reduced orphan risk for a given block size.
(Prop. 2) Continued existence of a transaction fee market.
(Prop. 3) Increased proof-of-work security from fee revenue.
(Prop. 4) Improved double-spend security for unconfirmed transactions.

Although these are obviously desirable properties, it is important to note that they may fail to emerge if the assumptions listed in Section 1 are not realistic.

For example, to show how subchains reduce orphan risk (Section 4), we assumed a simple linear model for the propagation time of block information between miners. In reality, propagation times have a stochastic element and also certain miners will be better connected to the network hash power than others.[40, 41] Despite these weaknesses, the model used captures the critical fact that communicating *more* information generally takes *longer*, and so we believe the results developed in Section 4 will apply, at least qualitatively, to a wide range of practical scenarios.

In Section 5 and 6, we assumed the market for block space was perfectly competitive when analyzing how large a rational miner would make his Δ-block. In practice, however, miners often use the default mining settings and thus leave profit opportunities on the table. In the future, we believe miners will spend more time optimizing these settings in order to increase their profit margins, as the mining industry becomes more competitive. The perfect-competition assumption will thus become increasingly accurate.

We also assumed that the protocol-enforced block size limit (if one exists) was greater than the free-market equilibrium block sizes produced by miners. 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 does today, the marginal orphan 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 reduced. The double-spend resistance of unconfirmed transactions would likewise suffer.

Our examination of double-spend security in Section 7 assumed that the majority of the hash power was default-compliant. If instead, the majority of the hash power were petty-compliant—willingly facilitating double-spend attacks when bribed—then we expect that the increased security for unconfirmed transactions delivered by the subchain technique would be lost (along with much of the technique's other benefits).

Other real-world effects were also overlooked in this paper. For example, subchains would 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.

Finally, 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. That said, the author believes that network-wide support for subchains would add significant transactional capacity and improve the user experience, helping to further advance the adoption of Bitcoin.

Acknowledgement

The author gratefully acknowledges the kind review and suggestions of Gavin Andresen, "awemany" and the three anonymous reviewers. He also acknowledges Gregory Maxwell for motivating him to investigate fee market dynamics in scenarios where information is pre-propagated, and he thanks the many thoughtful and intelligent individuals at the Bitcoin Forum (bitco.in) for hours of pleasant and productive discussion.

Notes and References

[1] "Total Number of Transactions" chart. *Blockchain.info* (13 December 2015) https://blockchain.info/charts/n-transactions-total

[2] "Median Transaction Confirmation Time (With Fee Only)" chart. *Blockchain.info* (13 December 2015) https://blockchain.info/charts/avg-confirmation-time

[3] "Scalability." *Bitcoin Wiki* (13 December 2015) https://en.bitcoin.it/wiki/Scalability

[4] Murdoch, S. J., Drimer, S., Anderson, R., Bond, M. "Chip and PIN is Broken." *2010 IEEE Symposium on Security and Privacy*, Oakland, California (16 May 2010) http://www.unibank.org/toposign/chip_and_pin_is_broken.pdf

[5] Rizun, P. R. "A Transaction Fee Market Exists Without a Block Size Limit." No Publisher (2015) https://www.bitcoinunlimited.info/resources/feemarket.pdf

[6] Stone, G. A. "An Examination of Bitcoin Network Throughput Via Analysis of Single Transaction Blocks." No Publisher (2015) http://www.bitcoinunlimited.info/1txn

[7] Rizun, P. R. "The marginal cost of adding another transaction to a block is nonzero: empirical evidence that bigger blocks are more likely to be orphaned." *Reddit* (16 July 2016) https://www.reddit.com/r/btc/comments/4t6guk/the_marginal_cost_of_adding_another_transaction/

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.40

[8] Barski, C., and Wilmer, C. *Bitcoin for the Befuddled*. San Francisco: No Starch Press (2014)

[9] "Bitcoin Network Capacity Analysis – Part 6: Data Propagation." *Tradeblock Blog* (23 June 2015) https://tradeblock.com/blog/bitcoin-network-capacity-analysis-part-6-data-propagation

[10] Decker C., Wattenhofer R. "Information Propagation in the Bitcoin Network." *13th IEEE International Conference on Peer-to-Peer Computing*, Trento, Italy, September 2013

[11] Croman, K., *et al*. "On Scaling Decentralized Blockchains." *Financial Cryptography and Data Security 2016*. http://fc16.ifca.ai/bitcoin/papers/CDE+16.pdf

[12] "Perfect competition." *Wikipedia*. https://en.wikipedia.org/wiki/Perfect_competition

[13] 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

[14] Pseudonymous ("rocks"). Comment in "Gold Collapsing. Bitcoin UP." *Bitcoin Forum.* (12 November 2015) https://bitco.in/forum/threads/gold-collapsing-bitcoin-up.16/page-99#post-3585

[15] Andresen, G. "[Bitcoin-development] Weak block thoughts…" *Bitcoin-development* (23 September 2015) http://lists.linuxfoundation.org/pipermail/bitcoin-dev/2015-September/011157.html

[16] Pseudonymous ("TierNolan"). "Decoupling transactions and POW." *Bitcointalk* (18 April 2013) https://bitcointalk.org/index.php?topic=179598.0

[17] Andresen, G., Comment in "Faster blocks vs bigger blocks." *Bitcointalk* (3 July 2014) https://bitcointalk.org/index.php?topic=673415.msg7658481#msg7658481

[18] Rosenbaum, K., Russell, R. "IBLT and Weak Block Propagation Performance." *Scaling Bitcoin Hong Kong* (6 December 2015)

[19] Maxwell, G. "[Bitcoin-development] Block Size Increase." *Bitoin-development* 7 May 2015 (accessed 13 December 2015) https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2015-May/007880.html

[20] BitFury Group. "Proof of Stake versus Proof of Work." No Publisher (13 September 2015) http://bitfury.com/content/4-white-papers-research/pos-vs-pow-1.0.2.pdf

[21] Andresen, G. "Back-of-the-envelope calculations for marginal cost of transactions." No Publisher (2013) https://gist.github.com/gavinandresen/5044482.

[22] Here we have made the assumption that the network hash rate is constant over the time scale at which new blocks are found. This assumption will likely no longer hold when fees are significant compared to the block reward, however, the underlying intuition regarding subchains effect on block races will be unchanged.

[23] If large miners are present, this equation overstates the orphan rate, due to the fact that a miner does not have to propagate his solved blocks to himself.

[24] The actual latency and propagation impedance are both likely smaller as the methodology used by Stone includes other effects such as the time to construct a new block candidate from mempool, and the methodology used by Tradeblock and Croman *et al*. measured propagation to nodes rather than to hash power.

[25] Pseudonymous ("awemany"). Comment in "Block Space as a Commodity." *Bitcoin Forum* (26 September 2015) https://bitco.in/forum/threads/block-space-as-a-commodity-a-

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.40

transaction-fee-market-exists-without-a-block-size-limit.58/page-4#post-1409

[26] Houy, N. "The Bitcoin Mining Game." *SSRN* (11 March 2014). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2407834

[27] "Law of supply." *Wikipedia*. https://en.wikipedia.org/wiki/Law_of_supply

[28] The derivative of this curve is a monotonically increasing function of Q via the law of supply, thus this curve is necessarily superlinear in Q.

[29] For more details, see Section 5 and Section 7 in the paper cited in endnote 5.

[30] For more details, see Section 4 in the paper cited in endnote 5.

[31] Pinna, D. "On the Nature of Miner Advantages in Uncapped Block Size Fee Markets." No Publisher (2015) http://www.scribd.com/doc/276849939/On-the-Nature-of-Miner-Advantages-in-Uncapped-Block-Size-Fee-Markets

[32] BitFury Group. "Incentive Mechanisms for Securing the Bitcoin Blockchain." No Publisher (2015) http://bitfury.com/content/4-white-papers-research/bitfury-incentive_mechanisms_for_securing_the_bitcoin_blockchain-1.pdf

[33] Carlsten, M., Kalodner, H., Narayanan, A. "Mind the Gap: Security Implications of the Evolution of Bitcoin Mining." *Scaling Bitcoin Montreal* (12 September 2015)

[34] The orphan rates for Litecoin and Dogecoin represent those experienced by the ProHashing pool and not the overall network orphan rates. https://prohashing.com/

[35] Todd, P. "Block Publication Incentives for Miners." No Publisher (29 June 2016). https://petertodd.org/2016/block-publication-incentives-for-miners

[36] Tschipper, P., "BUIP010: Xtreme Thinblocks." *Bitcoin Forum* (1 January 2016). https://bitco.in/forum/threads/buip010-passed-xtreme-thinblocks.774/

[37] Clifford, A., Rizun, P. R., Suisani, A., Stone, G. A., Tschipper, P. "Towards Massic On-chain Scaling: Block Propagation Results With Xthin. Part 5 of 5: Massive on-chain scaling begins with block sizes up to 20 MB." *Medium* (13 Jun 2016). https://medium.com/@peter_r/towards-massive-on-chain-scaling-block-propagation-results-with-xthin-5145c9648426

[38] Eyal, I., Gencer, A. E., Sirer, E. G., van Renesse, R. "Bitcoin-NG: A Scalable Blockchain Protocol." *arXiv* (7 October 2015). https://arxiv.org/abs/1510.02037

[39] McElrath, B. "Brading the Blockchain." Presentation at *Scaling Bitcoin Hong Kong* (7 December 2015). https://scalingbitcoin.org/hongkong2015/presentations/DAY2/2_breaking_the_chain_1_mcelrath.pdf

[40] Clifford, A., Rizun, P. R., Suisani, A., Stone, G. A., Tschipper, P. "Towards Massic On-chain Scaling: Block Propagation Results With Xthin. Part 3 of 5: Xthin blocks are less affected by the Great Firewall of China than standard blocks." *Medium* (4 Jun 2016). https://medium.com/@peter_r/towards-massive-on-chain-scaling-block-propagation-results-with-xthin-792a752c14c2

[41] Simply being a larger miner makes one "better connected to the network hash power" as one's connection with one's own hash power is usually very fast.

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.