NC-Max: Breaking the Security-Performance Tradeoff in Nakamoto Consensus

Ren Zhang*, Dingwei Zhang*, Quake Wang*, Shichen Wu†, Jan Xie* and Bart Preneel‡

*Nervos, †Shandong University, ‡imec-COSIC, KU Leuven

*{ren, dingwei, quake, j}@nervos.org, †shichenw@mail.sdu.edu.cn, ‡bart.preneel@esat.kuleuven.be

Abstract—**First implemented in Bitcoin, Nakamoto Consensus (NC) is the most influential consensus protocol in cryptocurrencies despite all the alternative protocols designed afterward. Nevertheless, NC is trapped by a security-performance tradeoff. While existing efforts mostly attempt to break this tradeoff via abandoning or adjusting NC's backbone protocol, we alternatively forward the relevance of the network layer. We identify and experimentally prove that the crux resides with the prolonged block propagation latency caused by not-yet-propagated transactions—fresh transactions. We thus present a two-step mechanism to eliminate fresh transactions and therefore remove the limits upon NC's performance imposed by its security demands, realizing NC's untapped potential to its maximum. Further, we introduce an accurate dynamic difficulty adjustment mechanism (DAM) to explore the real-time network condition and to adjust the protocol's throughput accordingly. Implementing the two-step mechanism and the DAM, we propose NC-Max, whose (1) security is analyzed, proving that it provides stronger resistance than NC against selfish mining and transaction withholding attacks, and (2) performance is evaluated, showing that it exhausts the full throughput supported by the network, and shortens the transaction confirmation latency by at least a factor of four compared to NC without compromising security.**

Index Terms—**Nakamoto consensus, throughput, selfish mining**

I. INTRODUCTION

Implementing *Nakamoto Consensus* (NC), Bitcoin [47], the most popular digital currency, allows all network participants to reach agreement on a chain of *blocks* containing confirmed transactions, and reignited the now well-known blockchain technology. In NC, *miners*—a special kind of network participants—compete for *block rewards* by solving a cryptographic puzzle generated from the latest block in the blockchain and a group of new transactions. A valid solution to the puzzle allows the miner to broadcast a new block packing these transactions, extending the blockchain. At the core of NC, a *backbone protocol* (1) periodically adjusts the puzzle difficulty via a *difficulty adjustment mechanism* (DAM), and (2) guides the miners to choose the same *main chain* when more than one block extends the same predecessor block.

As disruptive as Bitcoin is, its application is limited by its low throughput and long transaction confirmation latency, demanding further technological advances. Such a demand has been answered enthusiastically by both academia and the now hundred-billion-dollar industry of cryptocurrencies [2], where the legitimacy of new consensus protocols and new cryptocurrencies mostly resides in radical innovations and, particularly, outperforming NC. Consequently, a considerable number of new consensus protocols emerge, hoping to overcome NC's limitations by abandoning its backbone protocol. However, all of these new designs—represented by *proof-of-stake* (PoS) [19], [20], [31] and *blockDAG* protocols [4], [7], [44], [65], [66]—introduce hard-to-solve challenges in their security or functionality, as they forgo NC's simplicity. Specifically, PoS protocols, which select participants to compose blocks based on their possession of some scarce resources, demand additional security assumptions and protection mechanisms to prevent attackers from generating conflicting histories. These assumptions, however, are often difficult to meet [6], [56], and new attack vectors emerge even when the protection mechanisms are in place [11], [27], [37], [50]. BlockDAG protocols, which replace NC's linear blockchain structure with a direct acyclic graph of blocks, either abandon the global order of transactions [65], therefore limiting the smart contract functionality, or do not specify their transaction fee distribution [4], [44], [66] or DAM [4], [7], [66], rendering a complete security analysis infeasible. Due to these limitations and uncertainty brought by such radical innovations, NC and its variants remain the foundation of most of the leading cryptocurrencies such as Bitcoin, Litecoin [1], Ethereum [13], Bitcoin Cash [8] and Zcash [63]. Also built on NC are some influential protocol designs, represented by Fruitchains [54] and Bitcoin-NG [25].

Nevertheless, a key challenge confronting these NC-based designs is to improve NC's performance without compromising security, due to a well-known tradeoff rooted in its security and performance's conflicting requirements on the block size and the block interval—its security demands small blocks and long block intervals, while its performance demands larger blocks and shorter intervals [18], [67]. To break this tradeoff, a great deal of endeavors has been directed to adjusting NC's backbone protocol [42], [60], [67]. However, it has been shown [39], [76] that such adjustments often complicate the protocol and thus also lead to new attack vectors.

While it is a general belief that the security-performance tradeoff is coupled with NC's backbone, we, in this paper, alternatively forward the importance of *the network layer* in breaking the tradeoff. We identify and experimentally prove that the tradeoff resides in the network layer via the existence of *fresh transactions*—transactions that have not finished propagating to the network—contained in a block. Therefore, to eliminate fresh transactions, we focus on the network layer and propose a two-step mechanism, ensuring all

transactions are *not* fresh when their full content is embedded in the blockchain. This mechanism, therefore, removes the limits on the block size & interval placed by the security demands. Meanwhile, to further strengthen security, we introduce a DAM in which *selfish mining*, the most influential attack against NC [26], is not profitable. Our DAM also dynamically exploits the lower limit on the block interval, reducing NC's transaction confirmation latency. Integrating these two mechanisms—two-step transaction confirmation and our DAM, we propose a new consensus protocol NC-Max, which not only maintains the same level of security as NC, but also achieves the full throughput supported by the network and reduces the transaction confirmation latency by at least a factor of four. Our contributions include:

Breaking NC's Security-Performance Tradeoff. We identify and remove the bottleneck of fresh transactions in this tradeoff. Through experiments deployed on, or with data collected from, the Bitcoin network, we unveil the detailed mechanism of how the block size & interval affect NC's security through the existence of fresh transactions. Specifically, a fresh transaction demands the nodes to request its content before forwarding the block to their peers. This extra request-and-reply round trip invalidates the *compact block* mechanism—Bitcoin's current block propagation acceleration technique [16]. The extended block propagation latency leads the protocol more vulnerable to various attacks [18], [21], [29], [67].

To remove this bottleneck, we introduce a two-step mechanism including *transaction proposal* and *commitment*, pipelining fresh transactions' synchronization and non-fresh transactions' confirmation, illustrated in Fig. 1. The resulting block-size-independent propagation latency reaches the lower limit permitted by the network, enabling more aggressive block size and block interval choices without compromising security.

Inheriting and Strengthening NC's Security. NC's security is carefully scrutinized [23], [26], [28], [29], [38], [53], [62], [72] thanks to its simple backbone protocol. We ensure that NC-Max inherits NC's backbone by re-coupling the required data of the two steps in an updated block structure so that no new message type is introduced. This approach simplifies both the security analysis—our protocol directly inherits NC's chain growth, chain quality, and common prefix properties [29], [53]—and the reward distribution, as there is no need to share the rewards among different types of blocks or to split the fees among blocks packing the same transactions.

We further strengthen the protocol's security by raising its resistance against *selfish mining*, which allows attackers to gain unfair block rewards by invalidating blocks mined by other miners [26]. We prove that, through incorporating non-main-chain blocks in the DAM, NC-Max renders selfish mining not profitable. The proof holds regardless of the attack strategy and duration.

Exploiting the Performance Limit. Our dynamic DAM maximizes the performance under real-time network conditions. Previous dynamic DAMs [14], [64] can only signal the adjusting direction—towards a longer or shorter interval. In
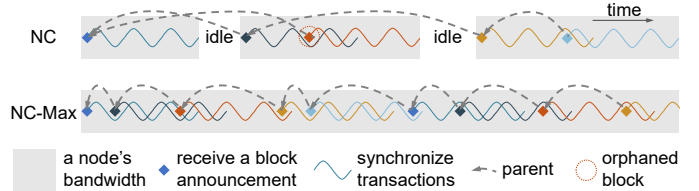

Fig. 1: NC-Max decouples transaction synchronization from confirmation to allow better bandwidth utilization and a shorter block interval without raising the orphan rate.

contrast, our DAM pinpoints the accurate expected interval matching the network condition through delicate modeling to explore the throughput limit without the possibly time-consuming trial-and-error procedure.

Our evaluation shows that NC-Max, under Bitcoin's current network condition and assuming a uniform bandwidth of 10 Mbps (1.25 MBps), achieves the nodes' full throughput of 2500 TPS. When anchoring the same sequence of block generation events, NC-Max's throughput outperforms recent high-throughput blockDAG designs. Anchoring an *orphan rate*—percentage of non-main-chain blocks—of 5%, NC-Max achieves a block interval of 5 seconds, whereas the interval must be over a minute in NC. Meanwhile, NC-Max reduces NC's more-than-six-minute transaction confirmation latency to 80 seconds, similar to that of Prism [4]—a blockDAG protocol characterized by its short latency.

Advocating a Broader View in Designing Consensus Protocols. Through this work, we aim to raise awareness of the importance of the network layer in designing protocols. Our protocol modification is built upon accurate identification and experimental verification of NC's network-layer bottleneck, which allows NC-Max to achieve better performance while preserving NC's security and flexibility. With NC-Max, we thus forward the configurational nature of a consensus protocol—that it shall be understood as a combination of the backbone protocol and the external rules—including those who concern the network layer, and the latter deserve no less attention for the former to exert its full potential and promise.

II. The Limitations of Blockchain Consensus Protocols

This section provides a brief overview of blockchain consensus protocols and their limitations, elaborating why we choose NC's backbone protocol as a base for improvements.

A. Nakamoto Consensus

NC helps all nodes agree on and order the set of confirmed transactions in a decentralized, pseudonymous way. There are two components in NC: *backbone protocol*, including its chain selection rules and DAM; *external rules*, including its transaction packing details, block validity rules, and reward distribution mechanism.

Backbone Protocol. Each block contains an 80-byte *header* in addition to the transactions. A block header includes (1)

the block's *height*—distance from the hard-coded *genesis block*, (2) the hash value of the *parent*—the latest block in the blockchain, (3) the Merkle root of the transactions, (4) a timestamp, and (5) a nonce. Embedding the parent hash ensures that a miner chooses its parent before starting to mine. Based on this parent-child relationship, all blocks form a tree, and each root-to-leaf path in the tree is called a *chain*. To construct a valid block, miners work on finding the right nonce so that the block hash is smaller than a *target T*, which is computed by the last iteration of the DAM. Compliant miners publish blocks the moment they are found.

When more than one block extends the same parent, miners work on the *main chain* that is most computationally challenging to produce, which is sometimes inaccurately referred to as the *longest chain*. When several chains are of the same "length", miners choose the first chain they receive. We refer to this *forked* situation where miners work on different parents as a *block race*, an equal-length block race as a *tie*, and blocks of the same height as *competing blocks*. Blocks not in the main chain are *orphaned* and discarded by all miners.

To adjust the target T based on the network hash rate, a DAM is triggered after every *epoch* of M blockchain blocks:

$$T_{i+1} = \begin{cases} T_i \cdot \frac{1}{\tau}, & L_i < L_{\text{ideal}} \cdot \frac{1}{\tau} \\ T_i \cdot \tau, & L_i > L_{\text{ideal}} \cdot \tau \\ T_i \cdot \frac{L_i}{L_{\text{ideal}}}, & \text{otherwise} \end{cases} \quad , \quad (1)$$

where i is the epoch number, L_{ideal} is the ideal time duration of an epoch, τ is a dampening filter to prevent rapid changes of T, and L_i is the actual time duration of the last epoch, as reported in the blocks. In Bitcoin, $M = 2016$, $\tau = 4$, and L_{ideal} is two weeks so that the average block interval is ten minutes if the mining power remains constant.

External Rules. The transactions must not conflict with those in previous blocks of the same chain. The size of a block must not exceed a predefined *block size limit*. Miners are incentivized by two kinds of rewards. First, a *block reward* is allocated to the miner of every blockchain block via a special *coinbase transaction* in the block. Second, the value difference between the inputs and the outputs in a transaction is called the *transaction fee*, which goes to the miner who includes the transaction in the blockchain. We omit other details as they are not relevant to this study.

Security-Performance Tradeoff. NC's security relies on the expected block interval being significantly larger than the block propagation latency [29], to minimize the number of blocks mined during other block's propagation. Violating this condition leads to a high *orphan rate*—percentage of orphaned blocks—which lowers the adversarial mining power threshold to secretly generate a longer chain, downgrading the system's security threshold [67]. However, short block propagation latency demands a small block size limit, which, together with a long block interval, results in low throughput and long transaction confirmation latency.

Whether to raise the performance at the risk of downgrading security is the most heatedly-debated topic in the Bitcoin community in recent years [18]. We note that even if we abandon security, simple re-parameterization does not allow us to fully exploit the nodes' bandwidth to confirm transactions, because orphaned blocks do not contribute to the transaction confirmation yet still consume bandwidth to propagate [67].

B. Innovative Consensus Protocols: Into the Unknown

Inspired by the success of NC, a considerable number of consensus protocols have emerged. These efforts can be categorized into two groups: NC's chain-based variants—which adjusts the backbone—and innovative protocols—which abandons the backbone. Nevertheless, while the variants of NC introduce new attack vectors where they modify the backbone [76], innovative protocols all lead to new and unsolved challenges in their security, performance, or functionality. We now briefly introduce the limitations of existing innovative protocols, which are further divided into two approaches.

Proof-of-Possession Protocols. PoS (e.g., Algorand [31], Ouroboros Praos [20], and Snow White [19]) and *proof-of-space* (e.g., Spacemint [52]) protocols avoid the energy consumption in NC by selecting participants to compose blocks based on their possession of some scarce resources. Since the resources are not consumed during the block generation, extra security assumptions and protection mechanisms must be in place to prevent attackers from constructing multiple history versions. These systems usually rely on stronger-than-NC synchrony and online assumptions, or even trusted parties to checkpoint the blockchain, which lead to new attack vectors if these assumptions are not met. As an example of the protection mechanisms, Algorand demands that each block, regardless of its size, be accompanied by a 300 KB certificate, comprised of hundreds of digital signatures [31]. Broadcasting these signatures negatively affects performance. More discussions on these limitations can be found in [6], [11].

BlockDAG Protocols. SPECTRE [65], Meshcash [7], PHANTOM, GHOSTDAG [66], Prism [4], and Conflux [44] suggest that, rather than referring to a single parent, a block contains hashes to all blocks the miner has received satisfying certain conditions. By confirming transactions with blocks not necessarily in a chain, these protocols hope to achieve higher throughput than chain-based protocols.

BlockDAG protocols' actual throughput is yet to be quantified, as it is difficult to model how much bandwidth is wasted due to transactions embedded multiple times in simultaneous blocks [43]. This *duplicate-packing problem* further complicates the transaction fee distribution, whose mechanism is omitted in PHANTOM, GHOSTDAG, Prism, and Conflux. Moreover, Prism, with its three kinds of blocks, introduces potential incentive issues in block reward distribution and new attacks against its DAM, which are not accounted for in its design. At last, as the global order of transactions is not known when constructing a block, transaction validity can only be evaluated after the neighboring blockDAG topology is settled, resulting in a long confirmation delay. Alternatively, SPECTRE abandons the global order, rendering it incompatible with the most popular smart contract programming model [13].

C. NC's Backbone as a Base for Future Endeavors

Despite the emergence of numerous alternatives, NC's backbone protocol still has a threefold advantage compared to its alternatives. First, it is built on only a minimum set of security assumptions. Consequently, its security is carefully scrutinized and well-understood [23], [26], [28], [29], [38], [53], [62], [72]. Alternative protocols often open new attack vectors, either unintentionally [37], [76] or by relying on assumptions that are difficult to realize. Second, coupling the block producer election and transaction confirmation minimizes the consensus protocol's communication overhead. In contrast, alternative protocols often demand a non-negligible communication overhead to certify that certain nodes have witnessed a block or to transmit the same transactions multiple times. Third, NC's chain-based topology ensures that a transaction global order is determined at block generation, which (1) minimizes the transaction verification and confirmation latencies, and (2) supports all smart contract programming models [13], [55].

Therefore, we argue that NC's backbone protocol provides the most promising base for future protocol designs. Nevertheless, NC's security-performance tradeoff signals some fundamental limitations in its design. As prior attempts trying to break this tradeoff via modifying the backbone often introduce new attacks, a natural question arises: is it possible to break the tradeoff while keeping the backbone protocol intact? Before answering that question, we need to pinpoint the bottleneck in this tradeoff.

III. BOTTLENECK IN NC'S SECURITY-PERFORMANCE TRADEOFF

NC's security limits its performance as larger blocks and shorter intervals—i.e., better performance—cause longer block propagation latency and more frequent blocks, respectively; both approaches raise the percentage of blocks that overlap each other's propagation. More overlaps lead to more orphaned blocks, which weaken security. The Bitcoin developers are among the first to recognize the importance of short block propagation latency. They, therefore, introduced *compact blocks* (CB) in 2016 to reduce this latency. However, with CB implemented, fresh transactions become the obstacle to lower the block intervals and larger block sizes. This section first introduces CB, and then elaborates on how fresh transactions defer the block propagation and enable a transaction withholding attack, followed by experiments that confirm the identified bottleneck.

A. Compact Blocks

We briefly overview the protocol here and refer interested readers to [16] for the specification. CB reduces the block propagation latency by optimistically not transferring full transactions in a block by default. Each node supporting CB enables the *HB mode* with the last three peers that have sent blocks to the node. Whenever a new block is available at an HB peer, the peer constructs a CB by replacing each transaction with a 48-bit shortid, and adding two extra fields:

(1) a connection-specific 64-bit random salt and (2) a list of transactions for which the peer is certain that the receiving node is not aware of them—just the coinbase transaction in the current implementation. The second field is called *prefilled transactions*. A shortid is computed as siphash-2-4(txid,SHA-256(*header*||*salt*)), where siphash-2-4 and SHA-256 are two hash functions, txid is the transaction hash, and *header* contains the block's metadata. Replacing transactions of typically two to five hundred bytes each with their shortids can compress a 1 MB block into around 13 KB [46], which is significantly faster to transfer. The CB is then sent to the node. After receiving the CB, the node computes shortids for transactions in its unconfirmed transaction pool, matches them with those in the CB, and requests the missing transactions. Once the node receives these transactions, it constructs and forwards relevant CBs to its other peers, and starts verifying the newly-received transactions in the meantime.

B. Fresh Transactions and Transaction Withholding Attacks

According to the Bitcoin developers, after implementing CB, the main contributor to the block propagation latency is the extra round trip caused by the missing transactions, if there is one [46]. Not all nodes have stored these *fresh transactions* in their memory pools, as these transactions are usually broadcast only a few seconds before they are mined in a block. This extra-round-trip delay is magnified as it usually takes several hops for the block to reach the entire network.

Furthermore, resourceful miners can raise their revenues by deliberately packing transactions known only to themselves, which we termed *transaction withholding attacks*. The slower block propagation gives them more time to mine on their blocks before other miners have received them, hoping to orphan honest blocks mined during this period as they do not extend the longest chain, thus constituting a de facto selfish mining attack. As selfish mining's profitability raises superlinearly with the attacker's mining power, more resourceful miners have less incentive to accelerate their blocks' propagation [26], [46]. This attack differs from selfish mining in that the attacker does not delay his blocks further than their propagation latency.

As of 2020, the time of this writing, most Bitcoin blocks contain no fresh transactions, thanks to the ten-minute block interval. Since it takes only 15 seconds for a transaction to reach 90% of nodes [22], fresh transactions constitute only a small fraction of the transactions broadcast after the last block. This fraction increases if the block interval is shortened.

C. Confirming the Bottleneck

Fresh Transactions Affect the Block Propagation Latency. We first study the relationship between the number of fresh transactions, denoted as n_{fresh}, and the one-hop block propagation latency δ. The multi-hop latency is studied in Sect. VII.

The one-hop latency is further decomposed into two parts: $\delta = \delta_{\text{CB}} + \delta_{\text{prep}}$, where δ_{CB} denotes *the CB transmission latency* from a node's upstream peer to the node and δ_{prep} denotes *the block preparation latency*, i.e., the time between
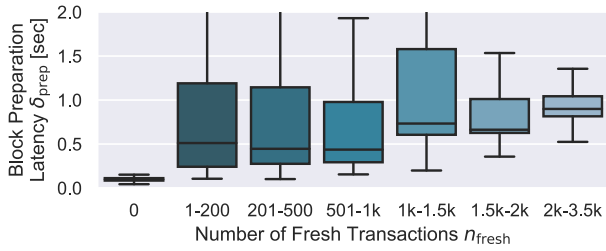
Fig. 2: The block preparation latency δ_{prep} is higher when a block contains fresh transactions, i.e., $n_{\mathrm{fresh}} > 0$. The ends of a box are the 25th and 75th percentiles of the data set; a horizontal line inside the box marks the median. The maximum latency is not displayed if it is over two seconds.

Fig. 3: The percentage of fresh transactions in a block p_{fresh} increases when the block interval $\overline{t_{\mathrm{in}}}$ decreases.

the node's receiving a CB and forwarding CBs to its peers. As CBs are small, δ_{CB} equals the latency between the upstream peer and the node, typically 1 to 50 ms [22], as upstream peers are usually a node's most well-connected peers. The extra round trip to request missing transactions is in δ_{prep}.

As δ_{CB} is relatively short and independent of n_{fresh}, we focus on the relation between n_{fresh} and δ_{prep}. We modified the Bitcoin client v0.17.1 to reject a fraction of new transactions to increase n_{fresh} when receiving a new block, and to log n_{fresh} and δ_{prep} for every block it receives. We deployed three instances of the modified client on IP addresses 47.251.4.119, 47.244.165.26, and 8.208.203.101, located in the US, Hong Kong, and the UK, respectively, and collected data from the Bitcoin network between May 10th and May 18th, 2019.

As shown in Fig. 2, when there is no fresh transaction, δ_{prep} is short and stable, typically around 100 ms. In this case, the latency is the time to reconstruct the block from the CB and to generate CBs for its peers. When $n_{\mathrm{fresh}} > 0$, even if the number is small, the median of δ_{prep} raises to over 500 ms, due to the request-and-reply round trip. The latency becomes even higher when there are more than 1000 fresh transactions, as larger messages take longer to transmit.

The Fraction of Fresh Transactions Increases When the Block Interval Decreases. Next, we study the relation between the expected block interval, denoted as $\overline{t_{\mathrm{in}}}$, and the fraction of fresh transactions, denoted as p_{fresh}.

We simulate the Bitcoin network with a higher transaction processing workload of 100 TPS. Our simulation setting is in Appendix D. A miner packs up to $100\overline{t_{\mathrm{in}}}$ transactions from its memory pool, prioritizing those with the oldest timestamps, as the block. "Packing from the oldest" results in a lower p_{fresh} than in reality, where miners pack from the highest fee-per-byte transaction. Each block's p_{fresh} is averaged over all the receiving nodes.

The results in Fig. 3 show very few fresh transactions in blocks when $\overline{t_{\mathrm{in}}} = 120$ seconds, which is consistent with Bitcoin's current situation. However, p_{fresh} increases when the block interval decreases, especially when it drops below the total transaction propagation latency.

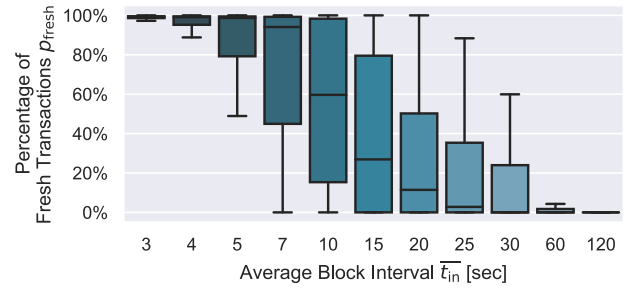Combining these results, we conclude that, when reducing the block interval, more fresh transactions appear in the blocks, which prolong the block propagation. Previous studies [21], [67] demonstrate that increasing the block size also raises this latency as more transactions take longer to synchronize, which will be further confirmed in Sect. VII. Longer block propagation latency directly results in a higher orphan rate, harming security. Next, we present our solution to this bottleneck.

IV. TWO-STEP CONFIRMATION

Fresh transactions invalidate the CB mechanism as they recouple the block propagation latency with the block size & interval. If we can ensure that the block propagation latency is short and block-size-independent, the security-performance tradeoff is broken: we can shorten the block interval to the lower limit permitted by the security demands, and raise the block size until the throughput exhausts the nodes' bandwidth. Moreover, there is no need to modify NC's backbone protocol. This goal is achieved via our two-step transaction confirmation mechanism, which introduces two adjustments to NC:

Incorporating Uncles. Miners are requested to refer to *uncles*—orphaned blocks in the same epoch—by embedding their hashes in the blocks. Our uncle's definition differs from that of Ethereum [13] in that our uncle's validity does not consider how far away the uncle and the nephew's first common ancestor is. The reward distribution regarding uncles (Sect. IV-D) is also different from that of Ethereum.

Prescribing a Transaction Proposal Step. A *transaction proposal zone* is added to each block, containing txpids—the first several bytes of transaction hashes—of some possibly fresh transactions. We consider these transactions *proposed* in this block. Transactions proposed in uncles are considered proposed in their nephews so that uncles contribute to the transaction proposal. The set of full transactions in a block is now called the *transaction commitment zone*. The proposed transactions—unlike their txpids—are not part of the block, therefore they do not affect the block's validity: a block may still be valid if some txpids refer to malformed or double-spending transactions, or the miner refuses to provide the full content of proposed transactions. The uncles' proposal zones are part of the block, and optionally part of the CBs if they have not been synchronized in the CB's connection. The most important rule about the proposal zone is that a transaction
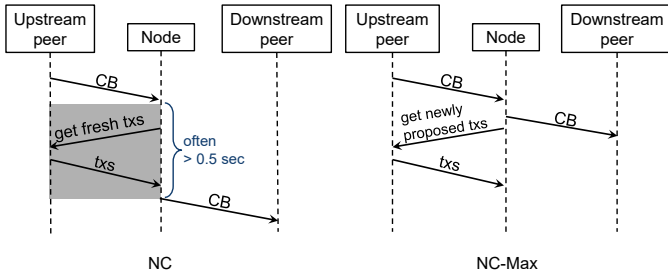
Fig. 4: Block propagation in NC and NC-Max. "Transactions" is abbreviated as "txs". On most occasions, our protocol allows nodes to forward CBs to their peers as soon as they received them, as all committed transactions are already synchronized.
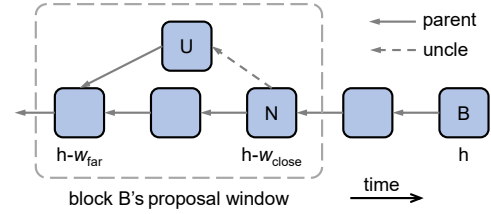

Fig. 5: Block B can only commit transactions proposed in its proposal window. In this example, $w_{\text{close}} = 2$, $w_{\text{far}} = 4$. If a transaction is only proposed in U and committed in B, its proposal fee goes to N's miner.

in the commitment zone of a block with height h must be proposed in a main chain block with height between $h - w_{\text{far}}$ and $h - w_{\text{close}}$, a range we termed *the proposal window*.

As proposed transactions do not affect a block's validity, a node forwards the CBs—including the full proposal zone—to its downstream peers as soon as it finishes reconstructing the commitment zone, whose transactions are already synchronized after receiving the blocks in the proposal window. The request to the node's upstream peer for missing proposed transactions in the latest block is sent in the meantime. Consequently, the round-trip time of requesting the missing transactions is removed from the critical path of block propagation, as shown in Fig. 4. Malicious miners may still conduct transaction withholding attacks by refusing to provide the full transactions they proposed and then hoping to commit these secret transactions in the future; however, the success rate and the damage of this attack are lower in NC-Max than in NC.

Next, we formally define the two steps and the block structure, and then introduce the new block propagation protocol and our reward distribution mechanism.

A. Definitions

Definition 1 (Proposal id). *A transaction's proposal id* txpid *is defined as the first l bits of the transaction hash* txid.

A txpid does not need to be globally unique as the 32-byte txid, as a txpid is used to identify a transaction among several neighboring blocks. Since we embed txpids in both blocks and CBs, sending only the truncated txids could reduce the bandwidth consumption. When multiple transactions share the same txpids, all of them are considered proposed. In practice, we can set l to be large enough so that the computational effort of finding a collision is non-trivial. For example, when $l = 144$, finding a txpid collision is roughly as difficult as mining a block in Bitcoin.

Definition 2 (Uncle). *A block B_1 is considered an* uncle *of another block B_2 if all of the following conditions are met: (1) they are in the same epoch, sharing the same difficulty; (2) $height(B_2) > height(B_1)$; (3) B_1's parent is already embedded in B_2's chain, either as a main chain block or as*

an uncle; and (4) B_2 is the first block in its chain to refer to B_1.

Condition (1) enables uncles to contribute to a more accurate hash rate estimation, which will be exploited in our DAM. A violation of (2) contradicts the longest-chain rule. Condition (3) is to ensure that two NC-Max instances with different genesis blocks do not accidentally recognize each other's blocks as uncles. An example is in Fig. 5.

Definition 3 (Transaction proposal). *A transaction is* proposed *at height h_p if its* txpid *is in the* proposal zones *of the main chain block with height h_p or of this block's uncles.*

The proposal zone facilitates transaction synchronization. The proposed transactions' validity does not affect the block's validity.

Definition 4 (Transaction commitment). *A non-coinbase transaction is* committed *at height h_c if all of the following conditions are met: (1) it is proposed at height h_p of the same chain, where $w_{\text{close}} \leq h_c - h_p \leq w_{\text{far}}$; (2) it is in the* commitment zone *of the main chain block with height h_c; and (3) it is not in conflict with any previously-committed transactions in the main chain. The coinbase transaction is* committed *at height h_c if it satisfies (2) and (3).*

A transaction is considered embedded in the blockchain when committed. We borrow the term *commit* from the literature for convenience, despite that in line with NC, NC-Max only offers probabilistic confirmation. Parameters w_{close} and w_{far} define the *proposal window*—the closest and farthest on-chain distance between a transaction's proposal and commitment, as shown in Fig. 5. Enforcing a proposal window guarantees that the "proposed transaction pool" fits in a node's memory. We suggest the proposal window $w = w_{\text{far}} - w_{\text{close}} + 1$ to be at least four to ensure liveness (Sect. VI-C). Although a longer window gives the miners more time to commit a transaction, a shorter window offers stronger resistance against transaction withholding attacks (Sect. VI-D).

We require w_{close} be large enough with a lower bound of two, so that w_{close} block intervals are long enough for newly-proposed transactions to finish propagation, and as small as possible to reduce the transaction confirmation latency. The selection of w_{close} will be further discussed in Sect. VII-D.

B. Block and Compact Block Structure

Data Structure. A block includes the following fields:

header	block metadata
commitment zone	full transactions
proposal zone	txpids
uncle headers	headers of the uncles
uncles' proposal zones	txpids in the uncles

The header contains the Merkle roots of the commitment zone, the proposal zone, and uncle headers, in addition to NC's header. Similar to NC, a CB replaces the commitment zone with the transactions' shortids, a salt, and a list of prefilled transactions. All other fields remain unchanged in the CB.

A *block size limit* is applied to the total size of the first four fields, to limit the data size to synchronize across the network along with each PoW solution. The uncles' proposal zones do not count in this limit as they are usually synchronized before the block is mined. The number of txpids in a proposal zone also has a hard-coded upper bound.

Parameter Recommendation. Two heuristic requirements can help practitioners choose the parameters. First, the upper bound on the number of txpids in a proposal zone should be no smaller than the maximum number of committed transactions, so that this bound is not the protocol's throughput bottleneck. Second, ideally, a CB should be no bigger than 80 KB. According to Croman et al. [18], messages no larger than 80 KB have similar propagation latency in the Bitcoin network in 2016; larger messages propagate slower as the nodes' bandwidth becomes the bottleneck. This number changes as the network condition improves. A typical set of parameters is in Appendix A.

C. Block Propagation Protocol

In line with [5], [26], [35], nodes should broadcast all blocks with valid PoWs, including orphans, as they may become uncles. Valid PoWs cannot be utilized to pollute the network, as constructing them is energy-consuming.

On most occasions, NC-Max's block propagation protocol (Alg. 1) removes the round trip of fresh transactions, as illustrated in Fig. 4, so that block propagation latency is constant regardless of how many transactions are proposed; when the round trip is inevitable, NC-Max ensures that it only lasts for one hop in the propagation and the additional latency is limited. This is achieved by the following three rules.

R1: non-blocking transaction query. As soon as the commitment zone is reconstructed, a node forwards the CBs to its downstream peers and queries the newly-proposed transactions from its upstream peers simultaneously (Line 16 in Alg. 1).

The block propagation will not be affected by these transaction queries as long as they are answered before the next w_{close}-th block is mined.

R2: missing transactions, now or never. If certain committed transactions are unknown to a CB receiver, the receiver queries the sender with a short timeout (Line 6 to Line 7). Failure to send these transactions in time leads to the receiver disconnecting the sender (Line 8 to Line 12). If the disconnected

---

Algorithm 1 Our Block Propagation Protocol

procedure OnReceiveCompactBlock(CB, $fromPeer$)
1: $freshCommitShortid = \emptyset$, $freshProposeTxpid = \emptyset$
2: add $CB.prefilledTx$ to $memoryPool$
3: **for all** shortid $\in CB.commitmentZone$ **do**
4:     **if** shortid.$tx \notin memoryPool$ **then**
5:         add shortid to $freshCommitShortid$
6: **if** $freshCommitShortid \neq \emptyset$ **then**
7:     request $freshCommitShortid.tx$ from $fromPeer$
8:     start timer t
9:     **if** $t = $ timeOut **and** no reply **then**
10:         drop the connection with $fromPeer$
11:         initiate the HB mode with another peer
12:         **return**   ▷ *missing tx in the commitment zone*
13: **for all** shortid $\in CB.commitmentZone$ **do**
14:     **if** shortid.$tx \notin CB.proposalWindow$ **then**
15:         **return** invalid block   ▷ *commit without proposing first*
16: Execute Line 17, 19 and 23 in parallel:
17: **for all** $tx \in freshCommitShortid.tx$ **do**
18:     verify tx's validity
19: **for all** $toPeer$ **do**   ▷ *forward the CB*
20:     construct CB_{toPeer} from CB
21:     add $freshCommitShortid.tx$ to $CB_{toPeer}.prefilledTx$
22:     send CB_{toPeer} to $toPeer$
23: **for all** txpid $\in CB.proposalZone$ **do**   ▷ *request new txs*
24:     **if** txpid.$tx \notin memoryPool$ **then**
25:         add txpid to $freshProposeTxpid$
26: **if** $freshProposeTxpid \neq \emptyset$ **then**
27:     request $freshProposeTxpid.tx$ from $fromPeer$
28:     add the replied transactions to $memoryPool$

---

sender was an outgoing connection, the receiver establishes a new connection to a random node. Moreover, the incomplete block will not be propagated further before receiving these transactions from another peer.

Proposed-but-not-received transactions are committed either (1) in a successful transaction withholding attack, or (2) when w_{close} consecutive blocks are mined before the transactions proposed in the first one are synchronized. If the upstream peer is honest, as in (2), a short timeout is adequate to transfer the missing transactions, as an honest upstream peer has already received them before sending the CBs. In the case of (1), the attacker cannot delay the first hop of the block propagation more than the timeout value without the block being discarded. In practice, we set the timeout to be 2 seconds.

R3: transaction push. If certain committed transactions are previously unknown to a CB sender, they will be embedded in the prefilled transaction list of the outgoing CBs (Line 19).

This rule removes the round trip if the sender and the receiver share the same list of proposed-but-not-broadcast transactions. In a transaction withholding attack, this rule ensures that the secret transactions are only queried in the first hop of the block's propagation, and then pushed directly to the receivers in subsequent hops.

D. Reward Allocation

A fixed block reward goes to every main chain block miner, which is calculated based on Eqn. (11) in Appendix B-C.

For each committed transaction, 60% of its transaction fee, denoted as the *commitment fee*, goes to the main chain block miner who commits it; while the other 40%, denoted as the *proposal fee*, goes to the earliest main chain block miner who proposes the transaction within the proposal window. This 60-to-40 fee allocation balances the miners' incentives to propose transactions and to extend the longest chain, in line with Bitcoin-NG [25]. Uncle miners do not get any fees.

NC-Max—unlike Ethereum—issues neither uncle rewards to compensate uncle miners, nor nephew rewards to incentivize miners to embed uncles in their blocks. This is because uncle and nephew rewards raise the selfish mining profit and lower the mining power threshold to perform the attack [51], [59].

Miners embed uncles for three kinds of benefits. First, embedding uncles allows more transactions to be proposed in a block, thus earning more proposal fees for the nephew miner. Second, when there are not enough transactions to propose in a nephew block, some transactions can be omitted if they are already proposed in the uncles, to further accelerate the nephew's propagation and lower its probability of being orphaned. Third, embedding uncles contributes to a more accurate estimation of the network hash rate, thus contributing to the system's selfish mining resistance (Sect. VI-E).

V. DYNAMIC DAM

Our two-step mechanism enables us to lower the expected block interval. The next challenge is to locate the interval that best utilizes the nodes' bandwidth without affecting security. To tackle this challenge, we introduce an accurate dynamic DAM that exploits the bandwidth utilization to the limit of the real-time network condition. Our goal is twofold: (**G1**) to render selfish mining unprofitable; (**G2**) to dynamically adjust the throughput based on the network's bandwidth and latency. Meanwhile, to maximize compatibility and attack resistance, our DAM needs to satisfy four constraints, in line with NC:

C1. All epochs have the same target duration L_{ideal}.
C2. The maximum block reward issued in an epoch $R(i)$ depends only on the epoch number i so that the rewards are distributed at a predetermined rate.
C3. The hash rate estimation of the last epoch does not change too fast, to prevent attackers from manipulating the DAM and forging a blockchain [5], even if some miners' network is temporarily controlled by the attacker.
C4. The expected block interval should abide by predetermined upper and lower bounds. The upper bound guarantees service availability; the lower bound guarantees that NC-Max does not generate more traffic than most nodes' capacity, thus ensuring decentralization.

To achieve **G1**, NC-Max incorporates all blocks, instead of only the main chain, in calculating the *hash rate estimation* of the last epoch, and then applies a dampening factor to the estimation so that the adjusted output conforms to **C3**. This output determines the computing efforts required in the next epoch for each reward unit. The effectiveness of this mechanism is proved in Sect. VI-E. To achieve **G2**, our DAM

targets a fixed orphan rate o_{ideal}, rather than a fixed block interval as in NC. As our two-step confirmation ensures a relatively stable block propagation process, we can solve the expected block interval matching the target orphan rate with the last epoch's duration, orphan rate, and the main chain block number. As the target epoch duration is fixed (**C1**), we can solve the next epoch's main chain block number, block reward, and difficulty target after applying several dampening factors and upper/lower bounds to safeguard **C2** and **C4**. We defer the detailed description of our DAM to Appendix B.

Combined with the two-step mechanism, targeting a fixed orphan rate allows us to pipeline the synchronization of previously-proposed transactions and the confirmation of recently-committed transactions, reducing NC's long idle time, as shown in Fig. 1.

VI. SECURITY ANALYSIS

Having introduced the core design, next, we analyze the protocol's security metrics, its resistance against transaction withholding and selfish mining attacks.

A. Threat Model

In line with prior studies [5], [26], [30], [35], [62], [74]–[76], we assume the total mining power remains unchanged throughout the attack, and the adversarial mining power share α stays below half of the total mining power. Under this environment, we assume one strong attacker who coordinates all adversarial mining power, which can cause more damage than multiple attackers acting separately [26]. The attacker may temporarily turn off some mining equipment, hoping to manipulate the DAM for higher overall profit. The attacker receives and broadcasts messages with zero propagation delay, and can arbitrarily reorder messages. However, the attacker cannot slow down honest blocks' propagation. Other miners abide by the protocol. We do not consider the effect of transaction fees [15], [45], [69], as it only makes up 1% of the miners' total rewards in Bitcoin [9]. Neither do we consider eclipse attacks [36], [48] or network partitions.

B. Analysis of the Backbone Protocol

The security of NC's backbone protocol is specified as the upper and lower bounds on the chain quality, chain growth, and common prefix metrics, first defined and proved by Garay et al. [29]. Subsequent studies [23], [28], [38], [53], [72] optimize these bounds and loosen the assumptions.

Our modifications to NC are limited to (1) the block validity rules—specified as the *content validation predicate* and the *input contribution function* in [29]—and (2) the DAM. NC's security proofs make only two assumptions on the block validity rules, both of which are satisfied by NC-Max: each block introduces enough entropy; a block mined by an honest miner is valid to the others. As for the DAM, the security proofs require that the block propagation latency be shorter than the block interval, which is guaranteed by our block interval lower bound (**C4**) and strengthened by the two-step confirmation. In sum, NC-Max is an instantiation of NC's

backbone protocol, and thus is compatible with its potential security updates.

C. Ledger Persistence and Liveness

Next, we analyze the robustness of NC-Max as a transaction ledger, specified as persistence and liveness properties in the literature. We omit the persistence proof as it is identical to that of NC's in [29]. The liveness proof, however, is not the same, as we modify the content validation predicate and the input contribution function.

In the following theorem, $\Pi_{\text{NC-Max}}$ denotes the NC-Max protocol, λ is a security parameter, w is the proposal window size $w_{\text{far}} - w_{\text{close}} + 1$, round is a small time unit, f is the probability of at least one honest party succeeds in finding a block in a round, ϵ is a small security margin that covers the probability of several very unfortunate events.

Theorem 1 (Liveness). *Under the Honest Majority Assumption, in $\Pi_{\text{NC-Max}}$, assuming $\lambda = w/2f$, if a valid transaction tx is given as input to all honest players continuously for at least $u = \lceil (4\lambda + w_{\text{far}}/f)/(1-\epsilon) \rceil$ consecutive rounds, then all honest parties will report $(\text{committed}, tx)$ more than w blocks from the end of the ledger, with probability at least $\text{P}_{\text{live}}^{\text{NC-Max}} = \text{P}^2$, where $\text{P} = 1 - e^{-\Omega(\epsilon^2 f\lambda)}$.*

Proof. First, as we prescribe $w \geq 4$, we have $\lambda \geq 2/f$, which enables us to invoke the chain growth, common prefix, and chain quality properties (Lemma 7, 14, and Theorem 16 in [29]) with the same λ. Invoking chain growth, after u rounds, any honest party's chain grows by at least $4\lambda f + w_{\text{far}} = 2w + w_{\text{far}}$ blocks. For each of these chains, invoking chain quality, with probability at least P, there is at least one honest block in the first w blocks that will propose tx. Let say it is the i-th block. Invoking chain quality again, with probability at least P, there is at least one honest block among the w-block sequence between block number $i + w_{\text{close}}$ and $i + w_{\text{far}}$ that would commit tx. Note that $i + w_{\text{far}} \leq w + w_{\text{far}}$. The last w blocks ensure a common prefix. $\square$

NC shares a similar liveness property with shorter waiting time $u = 4\lambda/(1-\epsilon)$ and larger liveness probability P. In practice, the extended waiting time is canceled by the shortened block interval. Although a larger w gives us larger $\text{P}_{\text{live}}^{\text{NC-Max}}$, a smaller w helps NC-Max to resist better against transaction withholding attacks, which are analyzed next.

D. Resistance Against Transaction Withholding attacks

Success Rate. We compare the fraction of *attacker blocks*—blocks mined by the attacker—that can be slowed down in NC and NC-Max, denoted as F_{slow}. We neglect the attack's outcome by assuming all the blocks are in the main chain. This simplification does not affect the comparison: the more attacker blocks delayed, the more honest blocks orphaned.

In NC, this attack can be performed with every attacker block, namely $F_{\text{slow}}^{\text{NC}} = \alpha$. In NC-Max, a block's propagation can only be delayed if it contains proposed-but-not-broadcast transactions. To trigger this case, the attacker needs two blocks not too far from each other in the chain: the older one to

propose these secret transactions, and the younger one to commit. Namely, for every attacker block with height h, it can commit secret transactions only if there is another attacker block between $h - w_{\text{close}}$ and $h - w_{\text{far}}$. The probability that there is such a block is $1 - (1-\alpha)^w$. Therefore, $F_{\text{slow}}^{\text{NC-Max}} = \alpha(1 - (1-\alpha)^w)$. We immediately have $F_{\text{slow}}^{\text{NC}} > F_{\text{slow}}^{\text{NC-Max}}$, namely, NC-Max offers stronger resistance than NC.

Simulation. We simulate both protocols to demonstrate NC-Max's resistance against this attack in a more realistic setting. Our simulation models how α, the proposal window size w, the maximum attack-incurred latency, and orphaned blocks affect the attacker's unfair percentage of main chain blocks. The results show that after incorporating these real-world factors, NC-Max enjoys a greater advantage over NC than that in the success rate analysis. The detailed model and our results are in Appendix C.

E. Selfish Mining Resistance

We prove the following theorem on the selfish miner's profitability:

Theorem 2 (Selfish mining profitability). *An attacker cannot gain more block reward per time unit than mining honestly, regardless of how many epochs the attack lasts or what strategy the attacker adopts, if the following conditions are met: (1) orphaned blocks are incorporated in estimating the last epoch's hash rate in the DAM, and (2) each epoch's total block reward R is fixed and all goes to non-orphaned blocks.*

Modeling the Attack. The attack lasts n epochs, with epoch number $1, 2, \cdots, n$. The total block reward per epoch R is constant throughout the attack. We add an epoch 0 before the attack launches and assume the attacker mines honestly with full mining power α in this epoch. Note that this is just to ensure that the difficulty is well-defined throughout the attack process. Whether or not an attack strategy generates continuous unfair profit is not affected by the initial difficulty. We assume all orphans are caused by the attack, namely, there is no naturally orphaned block. All epochs are long enough so that we can neglect deliberately excluded orphans and assume all orphans are included as uncles [29].

Notation. The time duration of epoch i is denoted as t_i. The compliant miners' mining power share is denoted as $c = 1 - \alpha$, and we immediately have $c > 0.5$. Within each epoch, the attacker distributes the mining power among honest mining, selfish mining, and idle. The attacker's honest mining power share in epoch i is denoted as b_i, which extends the main chain without orphaning any honest blocks. The selfish mining power share is denoted as a_i, which extends the main chain and orphans the same number of honest blocks. The remaining mining power share $\alpha - a_i - b_i$ is temporarily idle, hoping to lower the hash rate estimation $\hat{H}'_i$. By these definitions we have $b_0 = \alpha$, $a_0 = 0$, $a_i, b_i \geq 0$ and $a_i + b_i \leq 1 - c$. For the compliant miners, $c - a_i$ mining power share extends the main chain, and the rest a_i mines orphaned blocks.

The fraction of reward in epoch i that goes to the attacker is denoted as p_i^{reward}. If mining honestly, p_i^{reward} equals the

mining power share α, and all epochs have the same duration t_0. Therefore, the attacker's time-averaged reward is $\alpha \times \mathrm{R}/t_0$, where $\alpha \times \mathrm{R}$ is the attacker's total reward in an epoch. We normalize both R and t_0 to 1 for simplicity so that $\alpha \times \mathrm{R}/t_0 = \alpha = 1 - c$. Similarly, the attacker's time-averaged reward in an attack is $\sum_{i=1}^{n} p_i^{\text{reward}}/\sum_{i=1}^{n} t_i$, where p_i^{reward} satisfies

$$p_i^{\text{reward}} = \frac{a_i + b_i}{c + b_i} \ , \qquad (2)$$

where $c + b_i$ is the total mining power share in extending the main chain, $a_i + b_i$ is the attacker's share. According to our DAM, we have

$$t_i = \frac{c + a_{i-1} + b_{i-1}}{c + b_i} \ , \qquad (3)$$

where $c + a_{i-1} + b_{i-1}$ is the total mining power in the last epoch, estimated by our DAM by counting both main chain blocks and uncles. The main chain grows slower than epoch 0, i.e., $t_i > 1$, when $c + a_{i-1} + b_{i-1} > c + b_i$, and vice versa.

Proof. We need to prove that, regardless of how the attacker chooses n, a_i and b_i for $1 \leq i \leq n$,

$$\frac{\sum_{i=1}^{n} p_i^{\text{reward}}}{\sum_{i=1}^{n} t_i} \leq 1 - c \ ,$$

where the left is the time-averaged block reward with the attack, the right is that of honest mining. Let $\sigma_k = \sum_{i=1}^{k} p_i^{\text{reward}} - (1 - c) \sum_{i=1}^{k} t_i$, proving the inequality is equivalent to proving $\sigma_n \leq 0$. Let $d_i = a_i + b_i$, we will prove a stronger result of $\sigma_n \leq c + d_n - 1$ by induction. First,

$$\sigma_1 = p_1^{\text{reward}} - (1 - c)t_1 = \frac{d_1 - (1 - c)(c + d_0)}{c + b_1}$$
$$= \frac{c + d_1 - 1}{c + b_1} \leq c + d_1 - 1 \ .$$

Next, assuming for $k < n$, $\sigma_k \leq c + d_k - 1$ holds. When $k = n$,

$$\sigma_n = \sigma_{n-1} + p_n^{\text{reward}} - (1 - c)t_n$$
$$\leq c + d_{n-1} - 1 + \frac{d_n}{c + b_n} - (1 - c)\frac{c + d_{n-1}}{c + b_n}$$
$$= c + d_{n-1} - 1 + \frac{d_n - (1 - c)(c + d_{n-1})}{c + b_n} \ ,$$

now we only need to prove

$$c + d_{n-1} - 1 + \frac{d_n - (1 - c)(c + d_{n-1})}{c + b_n} \leq c + d_n - 1$$
$$\Leftrightarrow \frac{(c + b_n)(d_{n-1} - d_n) + d_n - (1 - c)(c + d_{n-1})}{c + b_n} \leq 0$$
$$\Leftrightarrow (c + b_n - 1)(d_{n-1} - d_n) + c(c + d_{n-1} - 1) \leq 0 \ .$$

When $d_{n-1} \geq d_n$, the last inequality holds. Now we need to prove that it holds for $d_{n-1} < d_n$. Let $\epsilon_n = d_n - d_{n-1} > 0$, the last inequality becomes

$$-\epsilon_n(c + b_n - 1) + c(c + d_n - \epsilon_n - 1) \leq 0$$
$$\Leftrightarrow c(c + d_n - 1) - \epsilon_n(2c + b_n - 1) \leq 0 \ .$$

$\square$

As the validity of this "incorporating uncles" mechanism does not rely on NC-Max's two-step confirmation and dynamic block interval, it can be implemented independently in NC to strengthen its security.

At last, we note that although an attacker cannot gain higher time-averaged profit, he can still raise the *relative revenue*, i.e., the proportion of block rewards, by invalidating honest blocks. This limitation is inherited from NC's backbone protocol. We hope to fix this limitation in future work.

## VII. Simulation and Implementation

In this section, we experimentally confirm that NC-Max breaks the security-performance tradeoff and compare the performance of NC-Max with existing designs. First, we measure the block propagation latency and the orphan rate of NC and NC-Max under several transaction processing workloads and block intervals in an environment simulating the Bitcoin network. The results show that NC-Max has a stable low orphan rate independent of the throughput, therefore allows a shorter block interval with the same level of security.

Then we compare the throughput and the transaction confirmation latency of NC-Max with NC and blockDAG protocols. The results demonstrate that NC-Max enables the full utilization of the nodes' bandwidth, whereas all blockDAG protocols suffer from the duplicate-packing problem. Furthermore, the transaction confirmation latency is similar to that of Prism, a recent blockDAG protocol featured by its short latency.

A. Experimental Setup

Network and Mining. Our environment emulates the Bitcoin network's scale, connectivity, node latency, block and transaction propagation latency, and mining power distribution. Our data sources and sampling mechanisms are in Appendix D. The emulated network differs from the Bitcoin network in two ways: we prescribe (1) a uniform bandwidth of 10 Mbps (1.25 MBps) and (2) a stable network topology, which does not account for heterogeneous node resources and the joining and leaving of nodes. Yet NC's block propagation latency and its distribution in our simulation match well with the Bitcoin network, verifying the reasonableness of the environment.

Transaction Packing and Data Collection. We consider three transaction packing workloads: 100, 1000 and 2500 TPS, and eight average block intervals $\overline{t_{\text{in}}}$: 1 to 5, 10, 20 and 40 seconds. In the 100 TPS setting, 100 transactions are generated per second; each block packs up to $100\overline{t_{\text{in}}}$ transactions. Other settings follow similar constraints. The 2500 TPS setting exhausts the nodes' 10 Mbps bandwidth to synchronize transactions. The number of fresh transactions n_{fresh} is sampled from our second experiment in Sect. III-C with $\overline{t_{\text{in}}}$, the time since last block t_{in}, and the transaction packing workload as inputs. Note that the initial transaction propagation does not affect the block propagation as blocks have priority in propagation over transactions. A simulation is executed 40 times; each instantiates a new network topology and generates 200 blocks.
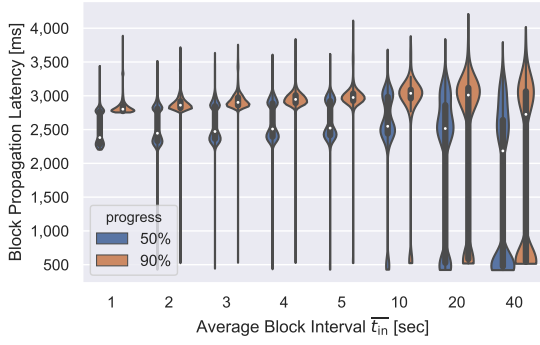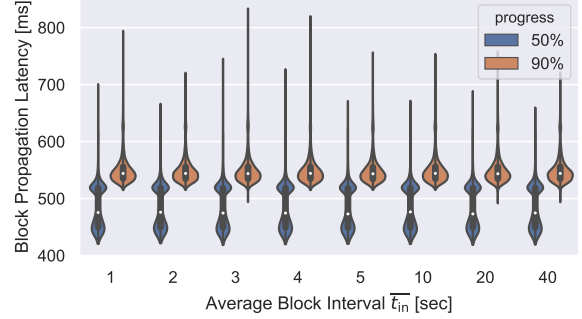
(a) NC, 100 TPS.
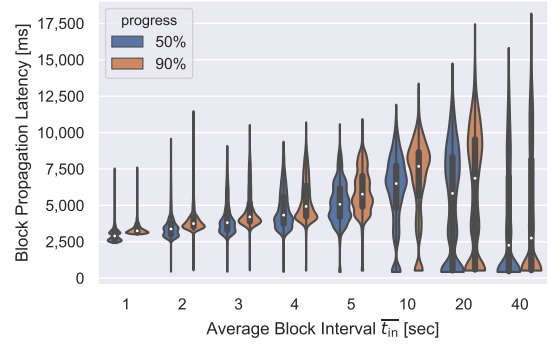
(b) NC, 1000 TPS.
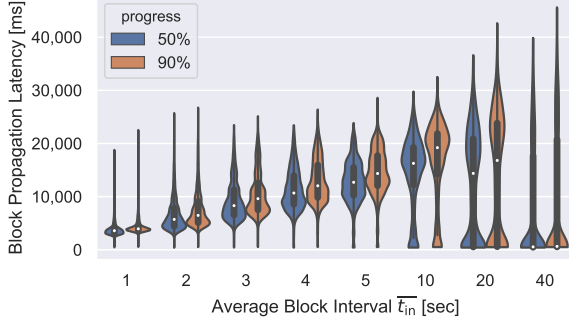
(c) NC, 2500 TPS.

(d) NC-Max, 100 TPS. For all data in this figure, $w_{\text{close}} = 2$.

Fig. 6: Blocks propagate faster in NC-Max than in NC. Latency distribution is displayed as the kernel density estimation [71]; wider value indicates higher density. The median is displayed as a white dot.
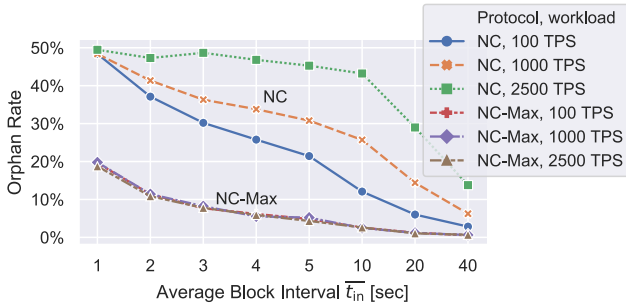

Fig. 7: Orphan rates. NC-Max with large-enough w_{close}.

B. Block Propagation Latency

Figure 6 shows the time distribution for blocks to propagate to 50% and 90% of nodes. We only display one setting for NC-Max as all settings follow similar patterns.

In the NC, $\overline{t_{\text{in}}} = 40$, 100 TPS setting, 50% of blocks contain no fresh transaction, which finish propagation within 600 ms; the other 50% take more than two seconds to finish propagation. These results match well with Bitcoin's current situation: when $\overline{t_{\text{in}}} = 600$, most blocks' propagation delay is within 600 ms, with a few exceptions over two seconds [22]. When $\overline{t_{\text{in}}}$ decreases, fewer and fewer blocks contain no fresh transaction. Also as $\overline{t_{\text{in}}}$ decreases, blocks with fresh transactions propagate slightly faster, because the block size decreases along with $\overline{t_{\text{in}}}$ in a fixed-TPS setting. However, the speedup is not proportional to the block size,

especially in the worst case, as a smaller $\overline{t_{\text{in}}}$ also means a higher fraction of fresh transactions. At last, blocks propagate slower as the workload increases, since the bandwidth, rather than the latency, becomes the bottleneck.

In NC-Max, when w_{close} is large enough, the latency is not affected by the block interval, and only marginally affected by the block size. In the 100 TPS setting, the 50% block propagation latency concentrates at two values: 450 ms and 520 ms, because sometimes it takes a block four hops to reach the 50th percentile, sometimes it takes five. Other settings have similar results, except for abiding by the bandwidth constraint.

C. Orphan Rate

In line with our DAM, an orphan rate here is computed as the number of orphaned blocks, divided by the number of main chain blocks. As displayed in Fig. 7, in NC, as the workload increases, the orphan rate deteriorates, since larger blocks take longer to propagate; whereas in NC-Max, the orphan rate is almost independent of the workload. Consequently, NC-Max reduces the block interval with the same orphan rate. For example, when $\overline{t_{\text{in}}} = 4$, the orphan rate is 6% in NC-Max; whereas in NC, the same orphan rate demands $\overline{t_{\text{in}}} = 20$ with 100 TPS, $\overline{t_{\text{in}}} = 40$ with 1000 TPS, $\overline{t_{\text{in}}} > 40$ with 2500 TPS.

D. Minimum Propose-Commit Distance

NC-Max's near-constant block propagation latency and low orphan rate require that newly-committed transactions have finished synchronization, which is ensured by the minimum

Fig. 8: The orphan rate increases when $w_{\text{close}} < w_{\text{close}}^{\min}$.

Fig. 9: In Nervos CKB, most epochs' duration is close to the four-hour target. Data collected between Nov. 16, 2019 and June 12, 2020.

propose-commit distance w_{close}. The minimum w_{close} to guarantee performance, denoted as $w_{\text{close}}^{\min}$, is a function of (1) $\overline{t_{\text{in}}}$, (2) the block size limit S, and (3) the network condition $\mathcal{Z}$, which encompasses the network's topology, the joining and leaving of nodes, and the nodes' heterogeneous resources.

We display $w_{\text{close}}^{\min}(\overline{t_{\text{in}}}, S, \mathcal{Z})$ and how the orphan rate raises when $w_{\text{close}} < w_{\text{close}}^{\min}$ in Fig. 8. For NC-Max, 100 TPS and all other omitted settings, $w_{\text{close}}^{\min} = 2$. The orphan rate is stable when $w_{\text{close}} \geq w_{\text{close}}^{\min}$; otherwise it is sensitive to w_{close} even when $\overline{t_{\text{in}}}$ is small: a change from 26 to 25 in the $\overline{t_{\text{in}}} = 1$, 2500 TPS setting doubles the orphan rate.

Due to the dynamic nature of $\mathcal{Z}$, we do not claim that our $w_{\text{close}}^{\min}$ values are directly applicable in practice. We suggest two mechanisms to ensure a large-enough w_{close}. First, to hard-code some lower bounds on w_{close} based on the expected block interval of the epoch. Second, to dynamically adjust w_{close} based on the block propagation information of the last epoch, in line with our DAM.

E. Throughput and Transaction Confirmation Latency

We simulate several blockDAG protocols and compare their throughput and transaction confirmation latency with NC and NC-Max. To ensure fairness, we apply the same sequence of block generation events to all protocols with each $\overline{t_{\text{in}}}$ and anchor the transaction processing workload at 2500 TPS. The memory pool size has a 300 MB limit—Bitcoin's peak size achieved in Dec. 2017 [12]; old transactions are dropped if it is full. We summarize the main results here and refer interested readers to the complete results in Appendix E.

Throughput. NC-Max outperforms all other protocols with all $\overline{t_{\text{in}}}$ values, reaching at least 2490 TPS when $\overline{t_{\text{in}}} < 40$. When $\overline{t_{\text{in}}} = 40$, some transactions are lost as sometimes the transactions generated between consecutive blocks exceed the pool size limit. If all miners pack randomly from their pool, blockDAG protocols perform only slightly worse than NC-Max—at least 2470 TPS when the memory pool is not full. However, their performance downgrades to that of NC if all miners share the same preference in transaction packing.

Transaction Confirmation Latency. We compare the latency of NC, NC-Max, and Prism [4]. Both NC-Max and Prism confirm transactions faster than NC, whose latency is at least four times that of NC-Max with the same orphan rate. The
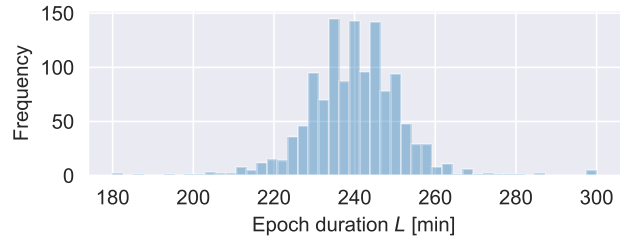
median latencies of Prism range from 50% to 90% of those of NC-Max when $\overline{t_{\text{in}}} < 5$, as all blocks in Prism contribute to transaction confirmation regardless of whether they are in the main chain. NC-Max outperforms Prism as $\overline{t_{\text{in}}}$ grows and orphaned blocks diminish. As an additional observation, blocks propagate faster are confirmed sooner in blockDAG protocols; whereas NC-Max confirms transactions at roughly even latency, providing more stable user experience.

VIII. IMPLEMENTATION

NC-Max is implemented in *Nervos CKB*, a public permissionless blockchain launched in Nov. 16, 2019, and has operated smoothly since then. The PoW puzzle of Nervos CKB is a dedicated hash function called EagleSong [3]. The project's consensus and peer-to-peer communication components consist of 14K lines of code (LOC). In addition, we implement Yamux [34] and libp2p secio [57] protocols as a separate library called *Tentacle* of 2.9K LOC to handle lower-level connection details. Nervos CKB[1] and Tentacle[2] are implemented in Rust and are open source.

The protocol parameters are instantiated as follows. The target epoch duration $L_{\text{ideal}} = 4$ hours. Coinbase transaction outputs, i.e., block rewards, can only be spent after four epochs. The lower and upper bounds on the expected block interval are 8 and 48 seconds, respectively. Each block embeds at most two uncles, which is enough given our target orphan rate $o_{\text{ideal}} = 2.5\%$. The maximum block size is 597 KB, which translates to around 1000 two-input-two-output committed transactions and a proposal zone of at most 1500 txids. The minimum and maximum propose-commit distances, i.e., w_{close} and w_{far}, are set to two and ten, respectively. Parameters of our DAM are in Appendix B.

The Nervos CKB network has grown considerably since its inception. The network hash rate, displayed in Fig. 10, has grown from 7.3×10^{13} to 1.1×10^{16} hash/sec. As of June 2020, the mining power is distributed among 12 mining pools with the largest one controlling roughly a third.

Nervos CKB demonstrates the effectiveness of our DAM from three aspects. First, as shown in Fig. 9, despite the 150-time hash rate growth in its first seven months, 93% of epochs' duration deviates no more than 20 minutes from

---

[1]https://github.com/nervosnetwork/ckb
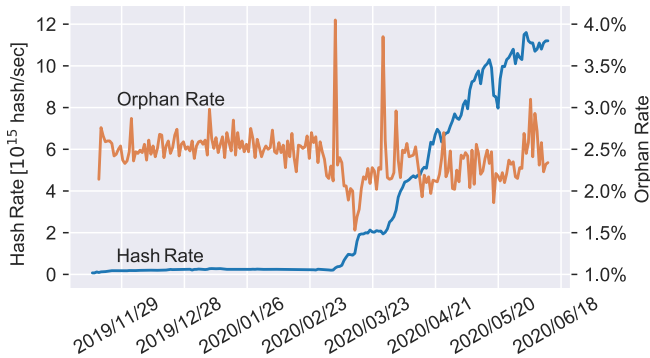[2]https://github.com/nervosnetwork/tentacle

Fig. 10: The network hash rate of Nervos CKB grows 150 times in its first seven months; the orphan rate stays close to the 2.5% target. The hash rate rocketed since early March as dedicated ASIC miners enter the market. Two spikes in the orphan rate correspond to two sharp increases in the hash rate. The initial orphan rate is omitted as it is close to zero.
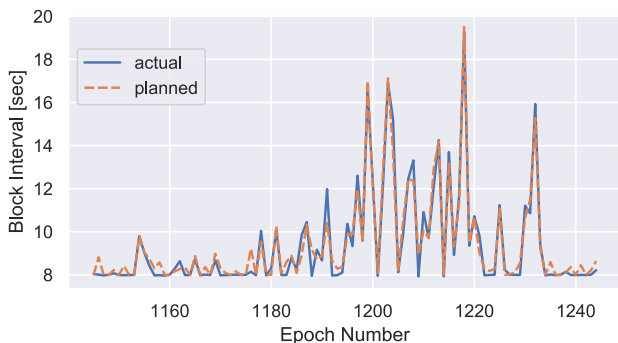

Fig. 11: The planned block interval outputted by our DAM matches well with the actual average block interval. We only display the last 100 epochs as the two lines become indistinguishable when including more epochs.

L_{ideal}. Second, as shown in Fig. 11, the measured average block interval in each epoch deviates only 3.7% on average from the planned expected interval outputted by our DAM; only 4.4% of epochs deviate more than 10%. These two aspects demonstrate the accuracy of our estimation of the total hash rate. Third, as displayed in Fig. 10, the orphan rate, targeting an ideal value of 2.5%, goes over 3% for only three days, proving that our adjustment on the difficulty, hence the block interval, stabilizes the orphan rate. Figure 11 also reveals that the planned block interval frequently hits the eight-second lower bound, indicating that the network is capable of processing more frequent blocks.

Meanwhile, the block propagation latency is satisfactory, which can be testified by the low orphan rate: Nervos CKB achieves a 2.5% long-term average orphan rate with a 7.5-second average block interval. The block interval is shorter than the lower bound as the hash rate constantly increases. In contrast, in Ethereum, which implements a variant of NC, the long-term average orphan rate is 10% [58]; the average block interval is 13 seconds [24]. Admittedly, these two systems are

not directly comparable given that Nervos CKB has a lower transaction processing workload. An in-depth analysis of the block propagation requires a well-connected network monitor, which we leave for future work.

IX. DISCUSSION

We now further elaborate on some possible concerns.

Remaining Performance Limit. While the block size limit can increase until the throughput exhausts most nodes' bandwidth, our protocol does not support arbitrarily low block intervals as it would lead to too high an orphan rate. For example, the orphan rate may undesirably exceed 10% and discourage some miners when $\overline{t_{\text{in}}}$ is one or two seconds. Accordingly, although NC-Max reduces the transaction confirmation latency of NC, a lower bound remains on this latency. Also, the correlation between the transaction processing workload and the confirmation latency remains, as transactions propagate slower when the nodes' bandwidth is exhausted.

The leading cause of orphaned blocks in NC-Max is the random nature of the mining process. Specifically, as block intervals follow an exponential distribution, it is impossible to enforce a lower bound on all of them. One approach to modify the distribution is to combine the current reversing-a-hash-function puzzle with a verifiable delay function [10].

Reasons to Embed Uncles. Another concern involves the miners' possible lack of financial incentives to embed uncles when there are not enough transactions to propose. Nevertheless, rational miners would still follow the protocol since, next to the proposal fees for embedding uncles which they are likely to receive under most situations, doing so would strengthen the system's security and, in particular, lower the possibility that the miners' blocks are orphaned in the long term. Moreover, not all honest behaviors need to be incentivized by fees. For example, Bitcoin relies on all miners not to set their clocks collectively faster to exhaust all remaining block rewards; no cryptocurrency provides any reward for nodes to forward blocks and transactions.

Pipelining vs. Concurrency. There is an emerging awareness in blockchain designs towards parallel block processing to exhaust the nodes' bandwidth. While parallel processing is mostly through *concurrency* as in blockDAG protocols, we highlight the potential of *pipelining*. Compared to concurrency protocols, which often involve duplicate transaction packing [43] and complicated algorithms to order blocks [44], [65], [66], a pipelining protocol, such as NC-Max, enjoys the advantage of simplicity, which leads to stronger security and functionality. Admittedly, pipelining protocols, including the recent Byzantine fault tolerance protocol Hotstuff [73], mandate an additional latency. In NC-Max, however, this propose-commit distance is to ensure that all nodes have received the transactions, which is necessary for all blockchains—including blockDAGs—before the transactions are considered confirmed. Explicitly mandating this latency strengthens the security, as demonstrated in our analysis of transaction withholding attacks.

In sum, as our pipelining approach already enables the full utilization of the nodes' bandwidth, we avoid introducing more complex rules—which often lead to more attack vectors and more limitations on functionality—to the consensus protocol.

X. RELATED WORK

We discuss some related work here in addition to Sect. II-B. We omit alternative designs that leverage more security assumptions than NC's and refer interested readers to Bano et al. [6] for these non-PoW and hybrid consensus protocols. Neither do we include off-chain [33] and sharding [70] protocols, where not all transactions are synchronized by all nodes. These protocols are compatible with NC-Max, and therefore can further increase a system's throughput.

Parallel Blocks. The GHOST protocol designed by Sompolinsky and Zohar [67] demands that during a block race, miners select the branch with the largest number of blocks, rather than the longest chain. GHOST resists better against double-spending than NC when the block interval is short so that natural forks happen frequently. However, Kiayias and Panagiotakos indicated that GHOST enables an attacker to slow down transaction confirmation via a *balancing attack* [40].

The Inclusive protocol proposed by Lewenberg et al. [43] also enables transaction confirmation by competing blocks. NC-Max inherits this key idea by allowing uncles to propose transactions while avoiding redundant transmission as each transaction is proposed with its txpid.

Decoupled Consensus. Bitcoin-NG by Eyal et al. [25] decouples NC's leader election and transaction serialization, allowing the nodes' full bandwidth to be utilized to confirm transactions. Unlike NC-Max, which shortens NC's transaction confirmation latency, this latency remains the same in Bitcoin-NG as in NC.

Novel DAM. Miller suggested a DAM targeting a one-to-one orphan-to-main-chain-block ratio, to minimize the block interval and to remove the block timestamp [64]. However, it is unclear how to distribute rewards at a constant speed without the timestamps. Moreover, Rosenfeld indicated that such a high orphan rate downgrades the attack difficulty [61].

The fact that selfish mining's profitability roots in the DAM was observed by Gervais et al. [30] and proved by Grunspan, and Pérez-Marco [32]. The latter study suggests incorporating uncles in the DAM to thwart the attack. Their proof on the new DAM's effectiveness does not cover the case where the attacker adapts to the modified mechanism and temporarily turns off some mining equipment, as indicated by Negy et al. [49]. Ethereum also incorporates uncles in its DAM [14].

Block Propagation Acceleration Techniques. The FIBRE network maintained by Corallo [17] deploys several high-speed nodes across the globe to help miners synchronize blocks the moment they are found. Bloxroute by Klarman et al. [41] is a network design that deploys high-speed nodes that start to relay blocks before receiving the full content. Both approaches are centrally coordinated. Although centralized systems can further accelerate the block propagation, the P2P acceleration techniques, including the one we proposed, serve as safety nets that are immune to single points of failure.

XI. CONCLUSION

While the current scholarly and engineering efforts to improve blockchains' performance concentrate on designing innovative consensus protocols, we alternatively highlight the importance of the network layer in the issue. By optimizing NC's block propagation mechanism on both the consensus and the network layer with a two-step mechanism, we break the throughput limit of NC without compromising security. We believe such an "incremental" approach, though not necessarily in line with the blockchain industry's eagerness for "revolutionary breakthroughs", is meaningful and inspirational. It reveals the significance of a long-overlooked aspect of protocol design—the network layer, and even more importantly, preserves the value and allows an easy transfer of our knowledge around NC—the knowledge that we, as scholars and engineers, collectively accumulated through the last decade.

REFERENCES

[1] "Litecoin," https://litecoin.org/.
[2] "Crypto-currency market capitalizations," 2020, https://coinmarketcap.com/charts/.
[3] Anonymized, "Our peer-reviewed hash function design," 2020.
[4] V. Bagaria, S. Kannan, D. Tse, G. Fanti, and P. Viswanath, "Prism: Deconstructing the blockchain to approach physical limits," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 585–602.
[5] L. Bahack, "Theoretical Bitcoin attacks with less than half of the computational power (draft)," *arXiv preprint arXiv:1312.7013*, 2013, https://arxiv.org/pdf/1312.7013.pdf.
[6] S. Bano, A. Sonnino, M. Al-Bassam, S. Azouvi, P. McCorry, S. Meiklejohn, and G. Danezis, "Consensus in the age of blockchains," *CoRR*, vol. abs/1711.03936, 2017, http://arxiv.org/abs/1711.03936.
[7] I. Bentov, P. Hubácek, T. Moran, and A. Nadler, "Tortoise and Hares consensus: the Meshcash framework for incentive-compatible, scalable cryptocurrencies," *IACR Cryptology ePrint Archive*, 2017, https://eprint.iacr.org/2017/300.pdf.
[8] bitcoincash.org, "Bitcoin cash," 2019, https://www.bitcoincash.org/.
[9] Blockchain, "Bitcoin block explorer," 2017, https://blockchain.info/.
[10] D. Boneh, J. Bonneau, B. Bünz, and B. Fisch, "Verifiable delay functions," in *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference*, ser. Lecture Notes in Computer Science, vol. 10991. Springer, 2018, pp. 757–788.
[11] J. Brown-Cohen, A. Narayanan, A. Psomas, and S. M. Weinberg, "Formal barriers to longest-chain proof-of-stake protocols," in *Proceedings of the 2019 ACM Conference on Economics and Computation*, ser. EC '19. ACM, 2019, pp. 459–473. [Online]. Available: http://doi.acm.org/10.1145/3328526.3329567
[12] BTC.com, "Unconfirmed transactions in Bitcoin," 2020, https://btc.com/stats/unconfirmed-tx.
[13] V. Buterin, "Ethereum: A next-generation smart contract and decentralized application platform," 2014, https://github.com/ethereum/wiki/wiki/White-Paper.
[14] ——, "Change difficulty adjustment to target mean block time including uncles," 2016, https://github.com/ethereum/EIPs/blob/master/EIPS/eip-100.md.
[15] M. Carlsten, H. Kalodner, S. M. Weinberg, and A. Narayanan, "On the instability of Bitcoin without the block reward," in *Proc. 2016 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '16. ACM, 2016, pp. 154–167, http://doi.acm.org/10.1145/2976749.2978408.
[16] M. Corallo, "Compact block relay," 2016, https://github.com/bitcoin/bips/blob/master/bip-0152.mediawiki.
[17] ——, "Public highly optimized fibre network," 2019, http://bitcoinfibre.org/public-network.html.

[18] K. Croman, C. Decker, I. Eyal, A. E. Gencer, A. Juels, A. Kosba, A. Miller, P. Saxena, E. Shi, E. G. Sirer *et al.*, "On scaling decentralized blockchains," in *Financial Cryptography and Data Security*, ser. LNCS, vol. 9604.   Springer, 2016, pp. 106–125.

[19] P. Daian, R. Pass, and E. Shi, "Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake," in *International Conference on Financial Cryptography and Data Security*. Springer, 2019, pp. 23–41.

[20] B. David, P. Gaži, A. Kiayias, and A. Russell, "Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*.   Springer, 2018, pp. 66–98.

[21] C. Decker and R. Wattenhofer, "Information propagation in the Bitcoin network," in *13th IEEE International Conference on Peer-to-Peer Computing (P2P)*.   IEEE, 2013.

[22] DNS Research Group, KASTEL @ KIT, "Bitcoin network monitor," 2019, https://dsn.tm.kit.edu/bitcoin/.

[23] T. Duong, A. Chepurnoy, and H.-S. Zhou, "Multi-mode cryptocurrency systems," in *Proceedings of the 2nd ACM Workshop on Blockchains, Cryptocurrencies, and Contracts*.   ACM, 2018, pp. 35–46.

[24] Ethstats, "Ethereum network status," 2020, https://ethstats.net/.

[25] I. Eyal, A. E. Gencer, E. G. Sirer, and R. V. Renesse, "Bitcoin-NG: A scalable blockchain protocol," in *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*. Santa Clara, CA: USENIX Association, 2016, pp. 45–59. [Online]. Available: https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/eyal

[26] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," in *Financial Cryptography and Data Security*.   Springer, 2014, pp. 436–454.

[27] P. Gai, A. Kiayias, and A. Russell, "Stake-bleeding attacks on proof-of-stake blockchains," in *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*, 2018, pp. 85–92.

[28] J. Garay, A. Kiayias, and N. Leonardos, "The Bitcoin backbone protocol with chains of variable difficulty," in *Advances in Cryptology – CRYPTO 2017*, ser. LNCS, vol. 10401.   Springer, 2017, pp. 291–323.

[29] J. A. Garay, A. Kiayias, and N. Leonardos, "The Bitcoin backbone protocol: Analysis and applications," in *EUROCRYPT*, 2015, pp. 281–310.

[30] A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdorf, and S. Capkun, "On the security and performance of proof of work blockchains," in *Proc. the 2016 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '16.   ACM, 2016, pp. 3–16, http://doi.acm.org/10.1145/2976749.2978341.

[31] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling Byzantine agreements for cryptocurrencies," in *Proceedings of the 26th Symposium on Operating Systems Principles*.   ACM, 2017, pp. 51–68.

[32] C. Grunspan and R. Pérez-Marco, "On profitability of selfish mining," *arXiv preprint arXiv:1805.08281*, 2018, https://arxiv.org/pdf/1805.08281.pdf.

[33] L. Gudgeon, P. Moreno-Sanchez, S. Roos, P. McCorry, and A. Gervais, "SoK: Off the chain transactions," Cryptology ePrint Archive, Report 2019/360, 2019, https://eprint.iacr.org/2019/360.

[34] Hashicorp, "Yamux (yet another multiplexer)," 2020, https://github.com/hashicorp/yamux.

[35] E. Heilman, "One weird trick to stop selfish miners: Fresh Bitcoins, a solution for the honest miner," Cryptology ePrint Archive, Report 2014/007, 2014, https://eprint.iacr.org/2014/007.

[36] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, "Eclipse attacks on Bitcoin's peer-to-peer network," in *24th USENIX Security Symposium (USENIX Security 15)*, 2015, pp. 129–144.

[37] S. Kanjalkar, J. Kuo, Y. Li, and A. Miller, "Short paper: I can't believe it's not stake! resource exhaustion attacks on PoS," in *Financial Cryptography and Data Security (FC) 2019*, ser. Lecture Notes in Computer Science, vol. 11598.   Springer, 2019, pp. 62–69.

[38] A. Kiayias, E. Koutsoupias, M. Kyropoulou, and Y. Tselekounis, "Blockchain mining games," in *Proceedings of the 2016 ACM Conference on Economics and Computation*.   ACM, 2016, pp. 365–382.

[39] A. Kiayias and G. Panagiotakos, "Speed-security tradeoffs in blockchain protocols," *IACR Cryptology ePrint Archive*, 2015, http://eprint.iacr.org/2015/1019.

[40] ——, "On trees, chains and fast transactions in the blockchain," *IACR Cryptology ePrint Archive*, 2016, https://eprint.iacr.org/2016/545.pdf.

[41] U. Klarman, S. Basu, A. Kuzmanovic, and E. G. Sirer, "bloxroute: A scalable trustless blockchain distribution network whitepaper," *IEEE Internet of Things Journal*, 2018.

[42] E. K. Kogias, P. Jovanovic, N. Gailly, I. Khoffi, L. Gasser, and B. Ford, "Enhancing Bitcoin security and performance with strong consistency via collective signing," in *Proc. 25th conference on USENIX Security Symposium*, 2016.

[43] Y. Lewenberg, Y. Sompolinsky, and A. Zohar, "Inclusive block chain protocols," in *Financial Cryptography and Data Security*, ser. LNCS, vol. 8975.   Springer, 2015, pp. 528–547.

[44] C. Li, P. Li, D. Zhou, Z. Yang, M. Wu, G. Yang, W. Xu, F. Long, and A. C.-C. Yao, "A decentralized blockchain with high throughput and fast confirmation," in *2020 USENIX Annual Technical Conference*, 2020, pp. 515–528.

[45] K. Liao and J. Katz, "Incentivizing blockchain forks via whale transactions," in *Financial Cryptography and Data Security*, ser. LNCS, vol. 10323.   Springer, 2017, pp. 264–279.

[46] G. Maxwell, "Advances in block propagation," 2017, https://www.youtube.com/watch?v=EHIuuKCm53o.

[47] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008, http://www.bitcoin.org/bitcoin.pdf.

[48] K. Nayak, S. Kumar, A. Miller, and E. Shi, "Stubborn mining: Generalizing selfish mining and combining with an eclipse attack," in *IEEE European Symposium on Security and Privacy (EuroS&P)*, IEEE.   IEEE, 2016, pp. 305–320.

[49] K. A. Negy, P. R. Rizun, and E. G. Sirer, "Selfish mining re-examined," in *Financial Cryptography and Data Security - 24th International Conference*, ser. Lecture Notes in Computer Science, vol. 12059.   Springer, 2020, pp. 61–78.

[50] M. Neuder, D. J. Moroz, R. Rao, and D. C. Parkes, "Defending against malicious reorgs in Tezos Proof-of-Stake," in *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, ser. AFT '20. Association for Computing Machinery, 2020, p. 4658.

[51] J. Niu and C. Feng, "Selfish mining in Ethereum," in *39th IEEE International Conference on Distributed Computing Systems, ICDCS 2019*.   IEEE, 2019, pp. 1306–1316.

[52] S. Park, A. Kwon, G. Fuchsbauer, P. Gaži, J. Alwen, and K. Pietrzak, "Spacemint: A cryptocurrency based on proofs of space," in *International Conference on Financial Cryptography and Data Security*. Springer, 2018, pp. 480–499.

[53] R. Pass, L. Seeman, and A. Shelat, "Analysis of the blockchain protocol in asynchronous networks," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*.   Springer, 2017, pp. 643–673.

[54] R. Pass and E. Shi, "Fruitchains: A fair blockchain," in *Proceedings of the ACM Symposium on Principles of Distributed Computing*, ser. PODC '17.   ACM, 2017, pp. 315–324, http://doi.acm.org/10.1145/3087801.3087809.

[55] A. Poelstra, "Scriptless scripts," 2017, https://download.wpsoftware.net/bitcoin/wizardry/mw-slides/2017-03-mit-bitcoin-expo/slides.pdf.

[56] A. Poelstra *et al.*, "Distributed consensus from proof of stake is impossible," *Self-published Paper*, 2014.

[57] Protocol Labs, "Libp2p secio: a secure transport module for go-libp2p," 2020, https://github.com/libp2p/go-libp2p-secio.

[58] YCharts Inc., "Ethereum uncle rate," 2019, https://ycharts.com/indicators/ethereum_uncle_rate.

[59] F. Ritz and A. Zugenmaier, "The impact of uncle rewards on selfish mining in Ethereum," in *2018 IEEE European Symposium on Security and Privacy Workshops (EuroS PW)*.   IEEE, April 2018, pp. 50–57.

[60] P. R. Rizun, "Subchains: A technique to scale Bitcoin and improve the user experience," *Ledger*, 2016, https://www.ledgerjournal.org/ojs/index.php/ledger/article/view/40.

[61] M. Rosenfeld, "Re: on the optimal difficulty setting for Bitcoin," 2012, https://bitcointalk.org/index.php?topic=98314.msg1075710#msg1075710.

[62] A. Sapirshtein, Y. Sompolinsky, and A. Zohar, "Optimal selfish mining strategies in Bitcoin," in *Financial Cryptography and Data Security*, ser. LNCS, vol. 9603.   Springer, 2016, pp. 515–532.

[63] E. B. Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, "Zerocash: Decentralized anonymous payments from bitcoin," in *2014 IEEE Symposium on Security and Privacy*.   IEEE, 2014, pp. 459–474.

[64] socrates1024, "on the optimal difficulty setting for Bitcoin," 2012, https://bitcointalk.org/index.php?topic=98314.msg1075701#msg1075701.

[65] Y. Sompolinsky, Y. Lewenberg, and A. Zohar, "SPECTRE: Serialization of proof-of-work events: Confirming transactions via recursive elections," 2016, https://eprint.iacr.org/2016/1159.pdf.

[66] Y. Sompolinsky, S. Wyborski, and A. Zohar, "PHANTOM and GHOSTDAG: A scalable generalization of Nakamoto Consensus," *IACR Cryptology ePrint Archive*, 2020, https://eprint.iacr.org/2018/104.pdf.

[67] Y. Sompolinsky and A. Zohar, "Secure high-rate transaction processing in Bitcoin," in *Financial Cryptography and Data Security*, ser. LNCS, vol. 8975. Springer, 2015, pp. 507–527.

[68] M. P. Stats., "Bitcoin mining pools," 2019.

[69] I. Tsabary and I. Eyal, "The gap game," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '18. ACM, 2018, pp. 713–728. [Online]. Available: http://doi.acm.org/10.1145/3243734.3243737

[70] G. Wang, Z. J. Shi, M. Nixon, and S. Han, "Sok: Sharding on blockchain," in *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, ser. AFT 19. New York, NY, USA: Association for Computing Machinery, 2019, p. 4161. [Online]. Available: https://doi.org/10.1145/3318041.3355457

[71] M. Waskom, "Kernel density estimation," 2019, https://seaborn.pydata.org/tutorial/distributions.html#distribution-tutorial.

[72] P. Wei, Q. Yuan, and Y. Zheng, "Security of the blockchain against long delay attack," in *Advances in Cryptology–ASIACRYPT 2018*, ser. LNCS, vol. 11274. Springer, 2018.

[73] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, "Hotstuff: BFT consensus with linearity and responsiveness," in *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, 2019, pp. 347–356.

[74] R. Zhang and B. Preneel, "On the necessity of a prescribed block validity consensus: Analyzing Bitcoin Unlimited mining protocol," in *13th International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*. ACM, Dec. 2017, pp. 108–119.

[75] ——, "Publish or Perish: A backward-compatible defense against selfish mining in Bitcoin," in *The Cryptographers' Track at the RSA Conference (CT-RSA)*, ser. LNCS, vol. 10159. Springer, Feb. 2017, pp. 277–292.

[76] ——, "Lay down the common metrics: Evaluating proof-of-work consensus protocols' security," in *40th IEEE Symposium on Security and Privacy (S&P)*. IEEE, May 2019, pp. 1190–1207.

APPENDIX A
PARAMETER EXAMPLES

Here is a typical set of parameters that satisfy the requirements mentioned in Sect. IV-B, assuming the block size limit is 1 MB. A header is 176 bytes, including 80 bytes of NC block header and three 32-byte Merkle roots. A proposal zone is 41 403 bytes at most, including 2300 18-byte txpids and a three-byte transaction count. For a block with one uncle, the header, the proposal zone, and the uncle header consume 41 755 bytes, which leaves 1 006 821 bytes in a block for the commitment zone, that can accommodate 2237 transactions of 450 bytes each. The corresponding CB includes 41 755 bytes for two block headers and the proposal zone, 13 422 bytes for shortids, 450 bytes for the coinbase transaction and less than 20 bytes of other metadata, in total less than 55 KB. The size becomes 95 KB if the uncle's proposal zone is included. Whether or not to include uncles' proposal zones in CBs depends on whether they are already synchronized in this connection. There is no need to propose transactions that are already proposed in the uncles, which reduces the size of CBs.

APPENDIX B
OUR DIFFICULTY ADJUSTMENT MECHANISM

A. Inputs and Outputs

Similar to NC, NC-Max's DAM is executed at the end of every epoch. It takes four inputs:

T_i last epoch's target

L_i last epoch's duration—the timestamp difference between epoch i and $i-1$'s last blocks

$C_{i,\mathrm{m}}$ last epoch's main chain block count

$C_{i,\mathrm{o}}$ last epoch's orphaned block count—the number of uncles embedded in epoch i's main chain

Among these inputs, T_i and $C_{i,\mathrm{m}}$ are decided by the last DAM iteration; L_i and $C_{i,\mathrm{o}}$ are measured after the epoch ends. The orphan rate o_i is calculated as $C_{i,\mathrm{o}}/C_{i,\mathrm{m}}$. We do not include $C_{i,\mathrm{o}}$ in its denominator to simplify the equations. As some orphans at the end of the epoch might be excluded from the main chain by an attack, o_i is a lower bound of the actual number. However, the proportion of deliberately excluded orphans is negligible as long as the epoch is long enough, as the difficulty of orphaning a chain grows exponentially with the chain length [29]. The algorithm outputs three values:

T_{i+1} next epoch's target

$C_{i+1,\mathrm{m}}$ next epoch's main chain block count

r_{i+1} next epoch's block reward

If the network hash rate and block propagation latency remain constant, o_{i+1} should reach the ideal value o_{ideal}, unless $C_{i+1,\mathrm{m}}$ equals its upper bound $C_{\mathrm{m}}^{\max}$ or its lower bound $C_{\mathrm{m}}^{\min}$. Epoch $i+1$ ends when it reaches $C_{i+1,\mathrm{m}}$ main chain blocks regardless of how many uncles are embedded.

B. Estimating Last Epoch's Hash Rate and Block Propagation Parameters

Adjusted Hash Rate Estimation. The adjusted hash rate estimation, denoted as $\hat{H}_i$, is computed by applying a dampening factor τ_1 to last epoch's *actual hash rate* $\hat{H}'_i$. The actual hash rate is calculated as follows:

$$\hat{H}'_i = \frac{\mathrm{HSpace}}{T_i} \cdot (C_{i,\mathrm{m}} + C_{i,\mathrm{o}})/L_i \ , \qquad (4)$$

where HSpace is the size of the entire hash space, e.g., 2^{256} in Bitcoin, HSpace/T_i is the expected number of hash operations to find a valid block, and $C_{i,\mathrm{m}} + C_{i,\mathrm{o}}$ is the total number of blocks in epoch i. $\hat{H}'_i$ is computed by dividing the expected total number of hash operations by the duration L_i.

Now we apply the dampening filter:

$$\hat{H}_i = \begin{cases} \hat{H}_{i-1} \cdot \frac{1}{\tau_1}, & \hat{H}'_i < \hat{H}_{i-1} \cdot \frac{1}{\tau_1} \\ \hat{H}_{i-1} \cdot \tau_1, & \hat{H}'_i > \hat{H}_{i-1} \cdot \tau_1 \\ \hat{H}'_i, & \text{otherwise} \end{cases},$$

where $\hat{H}_{i-1}$ denotes the adjusted hash rate estimation outputted by the last DAM iteration. The dampening factor ensures that the adjusted hash rate estimation does not change by more than a factor of τ_1—instantiated as 2 in Nervos CKB—between two consecutive epochs so that our DAM satisfies **C3**.

Modeling the Block Propagation. It is difficult, if not impossible, to model the detailed block propagation procedure, given that the network topology is unknown. Luckily, for our purpose, it suffices to describe the block propagation with two parameters, which will be used to compute $C_{i+1,\mathrm{m}}$ later. We learn some information on these parameters by computing the expected hash operations working on orphaned blocks in two different ways and then connect the expressions.

We assume all blocks follow a similar propagation model, in line with [18], [21]. In the last epoch, it takes d seconds for a block to propagate to the entire network, and during this process, the average fraction of mining power working on the block's parent is p. Therefore, during these d seconds, $\hat{H}'_i \times dp$ hash operations work on the latest block's parent, thus do not contribute to extending the blockchain. Consequently, in the last epoch, the total number of hashes that do not extend the blockchain is $\hat{H}'_i \times dp \times C_{i,\mathrm{m}}$.

On the other hand, the number of hash operations working on observed orphaned blocks is $\mathrm{HSpace}/T_i \times C_{i,\mathrm{o}}$. When the orphan rate is relatively low, we can ignore the rare event that more than two competing blocks are found at the same height. Therefore we have

$$\hat{H}'_i \times dp \times C_{i,\mathrm{m}} = \mathrm{HSpace}/T_i \times C_{i,\mathrm{o}} \ . \tag{5}$$

If we combine Eqn. (4) and (5), we can solve dp:

$$dp = \frac{\mathrm{HSpace}/T_i \times C_{i,\mathrm{o}}}{\hat{H}'_i \times C_{i,\mathrm{m}}} = \frac{o_i \times L_i}{(1+o_i)C_{i,\mathrm{m}}} \ , \tag{6}$$

C. Computing the Outputs

Main Chain Block Count. If the next epoch's block propagation situation is identical to the last epoch's, the value dp should remain unchanged. In order to achieve the ideal orphan rate o_{ideal} and **C1**—the ideal epoch duration L_{ideal}, following the same reasoning with Eqn. (6), we should have

$$dp = \frac{o_{\mathrm{ideal}} \times L_{\mathrm{ideal}}}{(1+o_{\mathrm{ideal}})C'_{i+1,\mathrm{m}}} \ , \tag{7}$$

where $C'_{i+1,\mathrm{m}}$ is the number of main chain blocks in the next epoch, if our only goal is to achieve o_{ideal} and L_{ideal}.

By combining Eqn. (6) and (7), when $o_i \neq 0$, we can solve $C'_{i+1,\mathrm{m}}$:

$$C'_{i+1,\mathrm{m}} = \frac{o_{\mathrm{ideal}}(1+o_i) \times L_{\mathrm{ideal}} \times C_{i,\mathrm{m}}}{o_i(1+o_{\mathrm{ideal}}) \times L_i} \ . \tag{8}$$

Now in order to achieve **C4**, we can apply the upper and lower bounds to $C'_{i+1,\mathrm{m}}$ and get $C_{i+1,\mathrm{m}}$:

$$C_{i+1,\mathrm{m}} = \begin{cases} \min\{C_{\mathrm{m}}^{\max}, \tau_2 C_{i,\mathrm{m}}\}, \\ \qquad o_i = 0 \text{ or } C'_{i+1,\mathrm{m}} > \min\{C_{\mathrm{m}}^{\max}, \tau_2 C_{i,\mathrm{m}}\} \\ \max\{C_{\mathrm{m}}^{\min}, C_{i,\mathrm{m}}/\tau_2\}, \\ \qquad C'_{i+1,\mathrm{m}} < \max\{C_{\mathrm{m}}^{\max}, C_{i,\mathrm{m}}/\tau_2\} \\ C'_{i+1,\mathrm{m}}, \qquad\qquad \text{otherwise} \end{cases} . \tag{9}$$

Equation (9) also covers the case of $o_i = 0$. When L_{ideal} is fixed (**C1**), a lower bound on $C_{i+1,\mathrm{m}}$ is equivalent to an

upper bound on the expected block interval $\overline{t_{\mathrm{in}}}$, and vice versa. The dampening factor τ_2, also instantiated as 2 in Nervos CKB, serves both **C3** and **C4**, preventing the main chain block number from changing too fast.

Target. To compute the target, we introduce an adjusted orphan rate estimation o'_{i+1}:

$$o'_{i+1} = \begin{cases} 0, & o_i = 0 \\ o_{\mathrm{ideal}}, & C_{i+1,\mathrm{m}} = C'_{i+1,\mathrm{m}} \\ 1/(\frac{(1+o_i)L_{\mathrm{ideal}}C_{i,\mathrm{m}}}{o_i L_i C_{i+1,\mathrm{m}}} - 1), & \text{otherwise} \end{cases} .$$

Using o'_{i+1} instead of o_{ideal} prevents some undesirable situations when $C_{i+1,\mathrm{m}}$ reaches its upper or lower bound. Now we can compute T_{i+1}:

$$T_{i+1} = \mathrm{HSpace}/\frac{\hat{H}_i \cdot L_{\mathrm{ideal}}}{(1+o'_{i+1}) \cdot C_{i+1,\mathrm{m}}} \ , \tag{10}$$

where $\hat{H}_i \cdot L_{\mathrm{ideal}}$ is the total number of hashes, $(1+o'_{i+1}) \cdot C_{i+1,\mathrm{m}}$ is the total number of blocks. The denominator in Eqn. (10) is the number of hashes required to find a block.

Note that if none of the edge cases is triggered, i.e., $\hat{H}_i = \hat{H}'_i$ and $C_{i+1,\mathrm{m}} = C'_{i+1,\mathrm{m}}$, we can combine Eqn. (4), (8), (10) and get $T'_{i+1} = T_i \times o_{\mathrm{ideal}}/o_i$. This result is consistent with our intuition. On the one hand, if o_i is larger than the ideal value o_{ideal}, the target lowers, increasing the difficulty of finding a block and raising the block interval if the hash rate is unchanged. Therefore, the orphan rate is lowered as it is more unlikely to find a block during another block's propagation. On the other hand, the target increases if the last epoch's orphan rate is lower than the ideal value, decreasing the block interval and raising the system's throughput.

Block Reward. Now we can compute the block reward:

$$r_{i+1} = \min \left\{ \frac{R(i+1)}{C_{i+1,\mathrm{m}}}, \frac{R(i+1)}{C'_{i+1,\mathrm{m}}} \cdot \frac{T'_{i+1}}{T_{i+1}} \right\} \ , \tag{11}$$

where the two cases differ only in the edge cases. The first case guarantees that the total reward issued in epoch $i+1$ will not exceed $R(i+1)$, thus ensuring **C2**. The second case ensures that an attacker that maliciously triggers some edge cases to make $T_{i+1} > T'_{i+1}$ cannot get more reward with the same mining power. When $o_i = 0$, only the first case applies.

In our security proof in Sect. VI-E, we use $r_{i+1} = R/C_{i+1,\mathrm{m}}$ and assume none of the edge cases is triggered. This is reasonable as: (1) when the mining power is stable, the dampening mechanism cannot be triggered without network-layer attacks; (2) causing $C_{i+1,\mathrm{m}} \neq C'_{i+1,\mathrm{m}}$ is not rational as it only lowers the attacker's block reward according to Eqn. (11).

Appendix C
Transaction Withholding Attack Simulations

Our simulation supports attackers with arbitrary mining power share α and arbitrary success rate of winning a tie, denoted as γ. To simplify the demonstration, we assume a
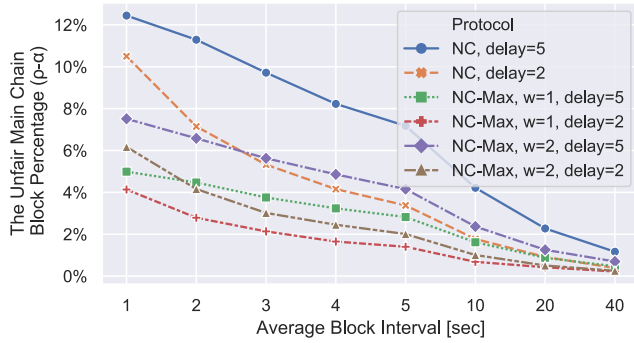
Fig. 12: NC-Max resists better against transaction withholding attacks than NC. The proposal window size $w = w_{\text{far}} - w_{\text{close}} + 1$; "delay" denotes t_{delay}, the maximum time an attacker can delay the propagation of its blocks.

strong attacker with $\alpha = 0.4$, who also wins all the ties—i.e., $\gamma = 1$, by immediately pushing his blocks to all the nodes after the competing blocks are mined. We omit natural orphans and assume that honest blocks are propagated immediately after they are mined, but not before the competing attacker block if there is one. The block interval follows an exponential distribution with expectation $\overline{t_{\text{in}}}$ as an input.

In both protocols, every "delayable" attacker block is delayed up to t_{delay}, during which the first honest block, if there is one, is orphaned. The attacker only delays the last block if he mines several blocks in a row. In NC, all attacker blocks are delayable. In NC-Max, the attacker delays his latest block only if there is another attacker block in the proposal window. The NC-Max simulation has an additional parameter $w = w_{\text{far}} - w_{\text{close}} + 1$. For every set of parameters, our simulation generates a chain of two million blocks and compute ρ, the proportion of main chain blocks mined by the attacker.

The results are displayed in Fig. 12. With the same t_{delay}, when $w = 1$ and 2, NC-Max reduces the damage $\rho - \alpha$ by roughly a half and a third, respectively. Although we compute the results where $t_{\text{delay}} = 2$ and 5 seconds for both protocols, in reality, t_{delay} is shorter in NC-Max than in NC, thanks to **R2** and **R3** introduced in Sect. IV-C.

APPENDIX D
SIMULATION ENVIRONMENT

Network. There are 6000 nodes, each establishes 8 random outgoing connections. The network latency between any two peers $\delta_{\text{a,b}}$ is sampled from data collected by a public Bitcoin crawler maintained by KIT [22]. All connections have the same bandwidth of 10 Mbps (1.25 MBps). The compact block (CB) transmission latency δ_{CB} equals $\delta_{\text{a,b}}$. Each time a node receives a CB, the block preparation latency δ_{prep} is a function

of n_{fresh}:

$$
\delta_{\text{prep}} = \begin{cases} \max\{n_1, 0\}, & n_{\text{fresh}} = 0 \\ \max\{1.33233 \times 10^{-4} n_{\text{fresh}} \\ \quad + 0.544959 + n_2, \\ \delta_{\text{a,b}} + 3.6 \times 10^{-4} n_{\text{fresh}}\}, & n_{\text{fresh}} > 0 \end{cases},
$$
(12)

where $\delta_{\text{a,b}} + 3.6 \times 10^{-4} n_{\text{fresh}}$ is the bandwidth constraint assuming each transaction is 450 bytes; n_1 and n_2 are random variables encompassing the variation of local-message-processing and network latencies: $n_1 \sim \mathcal{N}(0.101102, 0.0431702^2)$, $n_2 \sim \mathcal{N}(0, 0.420209^2)$, whose parameters have been learned via maximum likelihood estimation from the data we collected from the Bitcoin network (Sect. III-C). Parameters n_1, n_2, $\delta_{\text{a,b}}$, and δ_{prep} are in the unit of one second. Equation (12) is chosen based on our understanding of the block propagation: when $n_{\text{fresh}} = 0$, there is no round trip to query the fresh transactions; when $n_{\text{fresh}} > 0$, δ_{prep} grows linearly with n_{fresh}. We choose not to sample from the measurement data directly because that data misses certain n_{fresh} values. This delay also applies to NC-Max for querying transactions in the proposal zone. In both NC and NC-Max, the propagation of CBs is not affected by synchronizing fresh transactions; for each node, different blocks' δ_{prep} do not overlap each other to avoid violating the bandwidth constraint.

Mining. There are altogether 20 miners whose mining power distribution follows Bitcoin's on September 15, 2019 [68]. The top five miners control 17.42%, 17.39%, 16.63%, 13.51%, and 8.65% of mining power, respectively. The PoW is replaced with a scheduler that triggers block generation events at intervals following an exponential distribution with $\overline{t_{\text{in}}}$ as an input. An orphaned block emerges if it is mined before the latest block is propagated to its miner, therefore more than one block may be orphaned at the same height.

Transaction Generation and Propagation. A fixed number of transactions are generated per second, each from a random node as its initiator. Transactions are gossiped to the network from their initiators, with a fixed one-hop latency of 3.59 seconds so that the total propagation time—usually five hops—is similar to those of Bitcoin's 15 seconds [22]. A transaction is simulated as a dummy message with an ID and a timestamp. Each node maintains a memory pool of unconfirmed transactions.

APPENDIX E
COMPARISON WITH BLOCKDAG PROTOCOLS

A. Throughput

Modeling blockDAG Protocols. We model two types of blockDAG protocols: Prism [4] and the others [7], [44], [65], [66]. In Prism, each transaction has a *color*, possibly determined by its txid; all transactions in a *transaction block* are of the same color. Each miner maintains, for each color, a queue of unconfirmed transactions, and simultaneously mines on all queues. We simulate Prism with two, four, and six
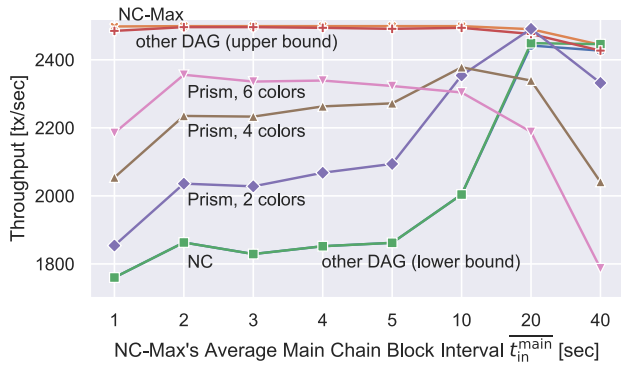
Fig. 13: NC-Max—in orange—outperforms other protocols in throughput. The lines of NC—in blue—and other DAG (lower bound)—in green—overlap except when $\overline{t_{\text{in}}^{\text{main}}} = 40$.
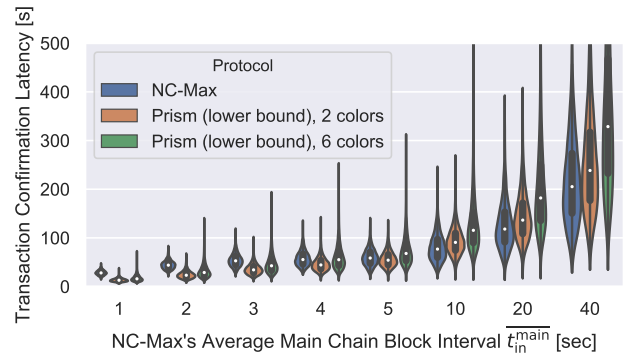

Fig. 14: NC-Max achieves similar transaction confirmation latency with Prism. The latency of Prism is simulated as lower bounds as we simplify its block-ordering logic. Values over 500 seconds are omitted.

colors. Other DAG protocols have no prescribed transaction inclusion rules, except that blocks do not include transactions in their predecessor blocks. For simplicity, we only implement two extreme transaction inclusion strategies: the lower bound and the upper bound. In the former setting, miners pack from the oldest transactions; in the latter, miners pack uniformly at random from their pools.

Additional Setup. To test whether NC-Max can exhaust the nodes' bandwidth, in each setting, we target a fixed main chain block interval $\overline{t_{\text{in}}^{\text{main}}}$ in NC-Max and adjust $\overline{t_{\text{in}}}$ accordingly. To ensure fairness, we apply the same sequence of block generation events to all protocols. Transactions are generated at an even speed of 2500 per second, capping the throughput as the physical limit. A block packs up to $2500 t_{\text{in}}^{\text{main}}$ transactions from the miner's pool. The pool size has a 300 MB limit—Bitcoin's peak size achieved in Dec. 2017 [12]; old transactions are dropped if it is full. In Prism, the pool size is evenly split among the colors.

Results. NC-Max reaches at least 2490 TPS when $\overline{t_{\text{in}}} < 40$, outperforming all other protocols (Fig. 13). When $\overline{t_{\text{in}}} = 40$, some transactions are lost as sometimes the transactions generated between consecutive blocks exceed the pool size limit.

Other DAG (upper bound) performs slightly worse than NC-Max, as a small fraction of transactions are included multiple times in *competing blocks*, wasting the processing capacity. Here we slightly abuse the term competing blocks to denote a group of blocks with no predecessor relation among them. Other DAG (lower bound) has almost the same throughput with NC, because when miners pack from the oldest transactions, competing blocks contain almost the same set of transactions unless the pool size is smaller than the block size—which only happens when $\overline{t_{\text{in}}^{\text{main}}} = 40$. In reality, the throughput of a blockDAG protocol should be between the lower and the upper bounds: to maximize their fees, miners prefer transactions with higher fees and but would randomize their selection to decrease the overlap with other miners' blocks [43].

Although Prism's throughput increases along with the num-

ber of colors, there are two caveats. First, more colors lead to longer transaction confirmation latency, as demonstrated in the next simulation. Second, the randomized nature of mining results in uneven processing speed among the queues, leading some queues to exceed their memory limits when some others are almost empty. In particular, when $\overline{t_{\text{in}}^{\text{main}}} > 10$, more colors lead to lower throughput. It is not trivial to dynamically allocate the memory among the colors without introducing new denial-of-service attack vectors.

We note that our results do not contradict blockDAG protocols' claims that they can exhaust the nodes' bandwidth. To do so, these protocols demand higher block frequency and larger memory pools than NC-Max.

B. Transaction Confirmation Latency

DAG protocols usually involve complicated transaction confirmation rules, therefore we only implement part of the transaction confirmation procedure of Prism to show that NC-Max achieves similar latency. We measure NC-Max's latency as $w_{\text{close}}^{\text{min}} + 6$ *main chain* block intervals after the transaction is broadcast: the first block to propose it, the $(w_{\text{close}}^{\text{min}} + 1)$-th block to commit it, and the last to settle it. In Prism, a transaction is confirmed in three steps: (1) a transaction block of the same color to embed it, which may take several block intervals; (2) a *proposer block* whose miner has received this transaction block to propose the latter; (3) several *voter blocks* to vote for this proposer block. Unlike NC-Max, there is no need for these blocks to be in the main chain in Prism. We consider a transaction settled when the sixth voter block is mined, regardless of the voter blocks' referring relationships. This simplification underestimates Prism's confirmation time, thus is in favor of Prism. We assume the same average block interval for each of the three kinds of blocks in Prism and blocks in NC-Max.

All three protocol instances displayed in Fig. 14 confirm transactions faster than NC, whose latency is at least four times that of NC-Max with the same orphan rate. Two Prism instances and NC-Max achieve similar latency. Specifically, the median latencies of two Prism instances are shorter than

NC-Max's when $\overline{t_{\text{in}}^{\text{main}}} < 10$ and 5, respectively, as all blocks in Prism contribute to transaction confirmation regardless of whether they are in the main chain. NC-Max outperforms them as $\overline{t_{\text{in}}}$ grows, because (1) orphan blocks diminish, and (2) there is no need to wait for a same-color transaction block. Moreover, the latency advantage of Prism, 2 colors comes at the price of lower throughput. Prism, 6 colors suffers from long worst-case confirmation latency—two to three times that of NC-Max, as it takes longer before a transaction meets a same-color block. As an additional observation, blocks propagate faster are confirmed sooner in blockDAG protocols; whereas NC-Max confirms transactions at roughly even latency, providing more stable user experience.