Modeling the Impact of Network Connectivity on Consensus Security of Proof-of-Work Blockchain

Yang Xiao*, Ning Zhang†, Wenjing Lou*, Y. Thomas Hou*
*Virginia Polytechnic Institute and State University, VA
†Washington University in St. Louis, MO

***Abstract*—Popularized by Bitcoin, proof-of-work (PoW) blockchain is one of the most widely deployed distributed consensus systems nowadays. Driven by incentives, PoW-based blockchain allows for democratized consensus making with correctness guarantee, as long as majority of the participants in the network are honest and rational. However, such elegant game theoretical security model falls apart when it is deployed on systems with potentially adversarial components and network conditions. For distributed consensus protocol used in blockchain, network plays a crucial role in the overall security of the system. A well-connected adversary with a communication advantage over honest nodes has a higher chance of winning blockchain forks and harvesting higher-than-usual mining revenue. Yet, evaluation of such communication advantage from a network perspective and its impact on blockchain consensus security have not received much attention.**

In this paper we fill this gap by assessing the impact of network connectivity on PoW blockchain consensus security via modeling analysis. Specifically, we perform analysis on two adversarial scenarios: 1) honest-but-potentially-colluding, 2) selfish mining. For each scenario, we evaluate the communication capability of nodes via network modeling and estimate the adversary's mining revenue and its impact on blockchain consensus security. Our analysis serves as a paradigm for future endeavors that seek to link blockchain security with network connectivity.

***Index Terms*—Blockchain, network modeling, consensus security**

I. Introduction

Decentralization is a foundational principle for blockchain technology and distributed ledger system. Envisioned by Nakamoto, the pseudonymous creator of Bitcoin [1], and later practitioners, blockchain essentially enables consensus making among mutually distrusted parties without relying on a central authority. A key premise that ensures the security of distributed consensus is that an adversary or a group of colluding adversaries do not control the majority of gross voting power in the consensus process, and in the case of proof-of-work (PoW) based blockchains, 50% of computing (or "mining") power of the entire network [2].

This honest-majority security premise comes under two assumptions. First, all nodes have the same communication capability, i.e. propagating information throughout the network equally fast. Second, during a blockchain fork race, wherein several blocks of the same height compete for a place in blockchain, all competitors have an equal chance of being the winner. However, in practice, the quality of connections often differ significantly among different network regions, as has been demonstrated by various measurements [3], [4], [5], [6]. Those residing in a highly connected cluster can disseminate blocks faster than those in a less connected region. This communication advantage translates into a higher chance of dominating a fork race, and has nontrivial consequence in the security of distributed consensus. As a result of this advantage at the network, the adversary will no longer require a 50% minimum of the gross mining power to succeed.

Following this intuition, various blockchain scaling proposals and security analyses [7], [8], [9] have identified the positive correlation between high blockchain fork rate and weak consensus security. These works generally adopt the *honest-but-potentially-colluding* threat model, in which any size of honest miners can join the collusion to compromise consensus security. Specifically, colluding miners can dominate fork races against the remaining honest miners and achieve unfair mining gains. As a result, the colluding miners may need less than 50% of mining power to break the consensus. However, these security analyses are largely qualitative and do not look into the impact of the actual network connection or information propagation dynamics.

The security impact of information propagation dynamics in Bitcoin was studied quantitatively at the macro level in [10]. It proposes a probabilistic model that estimates the average fork rate of Bitcoin blockchain based on the measurement of how an average block propagates in the network. The authors then regard fork rate as a security measure of the underlying blockchain network. However, this probabilistic model still assumes all miners have equal communication capability and equal chance of winning fork races, and does not consider the impact of heterogeneous network connectivity. It also does not provide a concrete case of how an adversary exploits the blockchain forks.

Another line of research focuses on adversarial strategies for selfish colluding parties [11], [12], [13], [14]. In selfish mining [11], an adversary with superior communication capability can achieve unfair mining gain by strategically withholding and releasing blocks. It proactively creates blockchain forks that nullify the efforts of honest nodes. Although these works take into account the difference in fork winning chance between the adversary and honest miners, their analyses treat the adversary's communication capability as a preexisting parameter (denoted by γ_{SM}) rather than deriving it from the actual network connectivity pattern. Additionally, how the

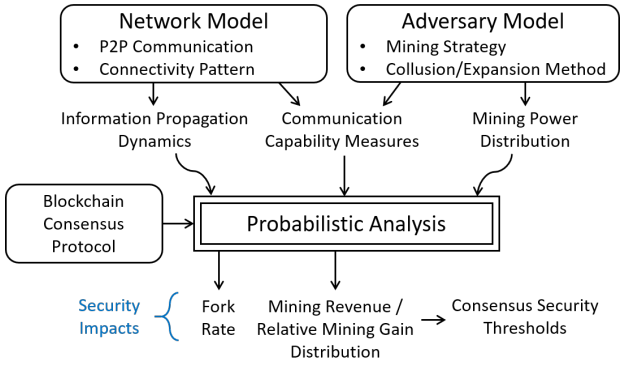

Fig. 1. Proposed analytical diagram.

expansion process of selfish mining pool in the network affects its communication capability and overall consensus security is also an important issue but overlooked in the literature.

Observing that both strategy and network capability are two key inter-connected factors for the security of distributed consensus, in this paper, we propose an analytical model to assess the impact of network connectivity on the security of PoW-based distributed consensus systems. The model captures network connectivity by a graph representation of the peer-to-peer network, and evaluates the communication capability of a miner node during a blockchain fork race. The communication capability measures, combined with the consensus protocol specification and other digests from the network model and adversary model, are then used to quantitatively estimate the security provisions of the whole system. An overview of our analytic diagram is illustrated in Figure 1.

Specifically, the main contributions of this paper include:

- We propose a novel analytical model that assesses the impact of network connectivity on consensus security of PoW blockchain through modeling and probabilistic analysis. The proposed analytical diagram also works for a general proof-of-X blockchain that relies on the longest-chain rule for consensus.
- For honest-but-potentially-colluding adversaries, our analysis treats every node as honest by default, and compute its mining revenue and relative mining gain based on its communication advantage over other nodes. With the distribution of mining revenues, we perform security analysis w.r.t. fork rate and the 50%-attack threshold.
- For selfish mining adversary, our analysis follows the classical setting of a selfish mining pool and treats the selfish mining pool as a consortium that expands among honest nodes. We provide a novel evaluation of the core communication capability metric of the pool, γ_{SM}, using a mining power-weighted betweenness centrality measure.
- We provide a thorough simulation experiment on PoW blockchain for each adversarial scenario. The simulation result validates our analysis.

II. Background

A. PoW Blockchain and Distributed Consensus

In public blockchain systems exemplified by Bitcoin, all networked miner nodes ("nodes" hereafter) work to curate a unified transaction history through distributed consensus. The transaction history is recorded in a chain of blocks in which every block contains a certain number of recently produced transactions. Every node seeks to generate the next block of the blockchain via a proof-of-work (PoW) process, namely, by finding an input to a cryptographic hash function that yields an output less than a target value. The input (i.e. the "proof") is attached in the block header. New blocks are disseminated immediately to the network via peer-to-peer gossiping. All nodes reach consensus on only one block at each blockchain height according to the "longest-chain rule": choosing the chain with the highest valid block. The generation of the next block should be aimed at prolonging the longest chain. Theoretically as long as the majority computing power is controlled by honest nodes, the longest chain shall always contain the most computation effort and thus the longest-chain rule ensures the security of distributed consensus. The above consensus scheme is also known as Nakamoto consensus, for its origin in Nakamoto's Bitcoin white paper [1].

In practical blockchain network, consensus security is complicated by blockchain forks. Blockchain fork is a scenario that multiple blocks of the same height are propagating in the network simultaneously. Under the assumption that all nodes are honest and follow the consensus rules, blockchain fork is caused by block propagation delays in that node j may generate a competing block before being aware of the existence of node i's block of the same height. To resolve blockchain fork, the longest-chain rule dictates that whichever fork branch gets appended with a new block should be chosen; blocks in other branches are then discarded. In the presence of forks, the honest-majority premise can still ensure consensus security, under an assumption that all competing blocks in a fork have an equal chance of being the winner [2].

B. Network Connectivity's Impact on Consensus Security

Due to heterogeneous connectivity of the underlying peer-to-peer network, the equal-chance fork winning assumption may not hold true. A well-connected node, say node i, tends to have superior communication capability that allows it to disseminate information faster than a less-connected node, say node j. If node i generates a new block, it takes a shorter time for node i to propagate this block across the whole network and thus the rest of the network has a lower chance of generating a competing block. If node j generates a competing block before node i's block reaches j, node i's communication advantage can still cause a larger share of the network to follow its block, which gives node i a higher chance of winning the fork eventually. As a result, in the long term well-connected nodes yield higher mining revenue than what would be expected from their share of computing power. This discrepancy between the long-term mining revenue and

the actual computing power of a node implies the possibility that a group of well-connected nodes with minority portion of computing power can harvest more than 50% of total mining revenue, which ultimately renders the honest-majority security premise vulnerable.

Besides exploiting naturally occurred forks, a well-connected adversary can achieve a significantly higher mining gain by proactively creating forks. Selfish mining [11] is one prominent example. Unlike an honest miner who publishes new block immediately after generation, a selfish mining attacker withholds newly generated blocks in a private chain, and strategically releases the private chain to the network whenever he sees his lead over the public chain decreases to a threshold. Consequently, the blockchain forks caused by the attacker's strategic private chain releases nullify the mining effort of honest nodes and create opportunities for the attacker to profit from its communication advantage. The detailed selfish mining strategy and the communication advantage parameter γ_{SM} will be discussed in Section V.

III. System Model

A. Network Model and Consensus Protocol

We consider a peer-to-peer network of N nodes represented by an undirected graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ and its adjacency matrix $\mathbf{A}$. $\mathbf{A}_{ij} = 1$ indicates node i, j share a peer relationship and can communicate in one hop. The PoW process and the longest-chain-rule-based consensus scheme are characterized as follows. To model the output randomness of the cryptographic hash function used for PoW, we assume each node i generates new blocks according to Poisson process of rate π_i per time unit δ. The block generation rate of the whole network is evaluated by the merged Poisson rate of $\pi = \sum_i \pi_i$. Note that our model does not adjust mining difficulty, as we are considering a fixed set of participants with fixed block generation rate.

Once a node i generates $block_i(h)$ of blockchain height h, it disseminates $block_i(h)$ throughout the network via peer-to-peer gossiping. Other nodes decide on the acceptance of $block_i(h)$ according to the longest-chain rule. That is, if another node k sees $block_i(h)$ while its local blockchain has already accepted $block_j(h)$ from node j ($j \neq i$), it declares a fork at height h and stops propagating $block_i(h)$. Conversely, if node l sees $block_i(h)$ before $block_j(h)$, it declares a fork at height h and stops propagating $block_j(h)$. Once the two competing blocks completely stop propagating and the network partitions into two factions each of which advocates one block, we call this situation a *fork stalemate*. And the two blocks are *partially propagated*. A fork stalemate can be resolved by a new block of height $h+1$ subscribing either i or j's block and being *fully propagated* in the network.

As for the finalization of blockchain, we consider the blockchain *canonized* by height h if a block of any origin $block_*(h)$ gets fully propagated in the network without encountering any competing block. We define the completion of $block_*(h)$'s propagation as a *canonization event*. Essentially, a canonization event at height h rejuvenates the block generation/competing process as if the past forks and competitions never happened. Figure 2 illustrates blockchain canonization and fork stalemate events. Note the canonization concept is different from probabilistic finality of Nakamoto Consensus [15], [16], which considers consensus security a probabilistic measure. We will use canonization events as embedding points to estimate the mining revenue of each node in the Section IV.

TABLE I
Summary of Notations

Network and Model Parameters	
$\mathbf{G}$	The graph representation of the node network.
N	Number of nodes in $\mathbf{G}$.
$\mathbf{A}$	The adjacency matrix of $\mathbf{G}$.
δ	Timeslot, also the time unit.
π	Block generation rate of the entire network (δ^{-1}).
π_i	Block generation rate of node i (δ^{-1}).
π_{SM}	Block generation rate of the selfish mining pool (δ^{-1})
Analyses	
$\mathbf{U}_i(t)$	Set of nodes unaware of node i's block at time t since its generation. $\lvert\mathbf{U}_i(t)\rvert$ is the cardinality of $\mathbf{U}_i(t)$.
$\lvert\mathbf{U}_i(t)\rvert_\pi$	Combined block generation rate of nodes in $\mathbf{U}_i(t)$.
$P_{NC,i}(t)$	Probability of the rest miners not proposing a competing block by time t of i's block's propagation.
$h(c)$	Blockchain height of the c^{th} canonization event
$\tau_{ij}(t)$	Minimum time for node i's block to reach j starting at time t from the generation of i's block.
$\omega_{i\succ j}(t)$	Node i's likelihood to win the fork race against node j if j publishes a competing block at time t from node i's block's generation. $\hat{\omega}_{i\succ j}(t)$ is an estimation.
γ_{SM}	Selfish mining pool's communication capability, i.e. the average fraction of honest mining power that will advocate the pool's block after it releases private chain.
MR_i	Mining revenue of node i as percentage of total canonized blocks.
RMG_i	Relative mining gain of node i. $RMG_i = \frac{MR_i - \pi_i/\pi}{\pi_i/\pi}$.
Security Metrics	
FR	Average fork rate of the whole network.
AT50	50%-attack threshold, i.e. minimum number of nodes whose combined mining revenue exceeds 50% of the total.
PRTH	Profitability threshold, i.e. pool size when selfish mining pool first achieves positive *RMG* during its expansion.

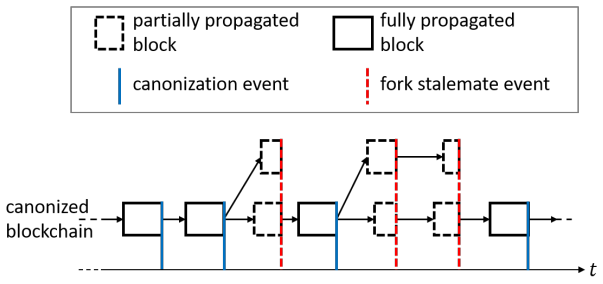

Fig. 2. Illustration of blockchain canonization and fork stalemate events. Width of a block denotes its propagation period.

B. Adversary Model

1) Honest-but-Potentially-Colluding: This adversarial scenario characterizes the practical case of the well-known 50% attack. That is, all nodes operate honestly by default, but the top miners can potentially collude so that their combined *mining revenue* (*MR*) exceeds 50% of the total. In our analysis, a well-connected node may obtain positive *relative mining gain* (*RMG*) and collude with other well-connected nodes. The mining revenue of a node can be viewed as its "enhanced mining power" in contrast to its actual computing power. In this scenario we are interested in the *50%-attack threshold* (*AT50*), i.e. the minimum fraction (computing power weighted) of the network whose aggregate mining gain exceeds 50% of the total.

2) Selfish Mining: This adversarial scenario assumes there are a pool of nodes in the network performing the selfish mining strategy as described in [11]. We treat the selfish mining pool as a consortium that expands among honest nodes and assumes their network connections. As the pool expands, it acquires the member nodes' computing power and external communication links. Under this scenario, *AT50* denotes the pool size when the pool first achieves 50% of total mining revenue during its expansion. We are also interested in the pool's *profitability threshold* (*PRTH*), which is defined as the pool size when the pool first achieves a positive *RMG*.

IV. Analysis on Honest Mining

In this section we calculate the impact of network connectivity on blockchain fork rate and mining gain distribution under the honest mining assumption. We then and discuss the security provision under the honest-but-potentially-colluding adversarial scenario.

A. Fork Rate

Define M_i as the event that node i is the first to generate the next block at an arbitrary moment of no outstanding blockchain fork. Denote the time for node i to find a block by random variable T_i. Then $T_i \sim exponential(\pi_i)$ and

$$P(M_i) = P\{T_i < T_j, \forall j \neq i\} = \frac{\pi_i}{\pi} \tag{1}$$

which can be conveniently derived from properties of Poisson processes. To facilitate the ensuing analysis, we consider the physical time slotted into basic time units of δ.

Let the moment when event M_i happens be time 0. Denote $\mathbf{U}_i(t)$ the set of nodes unaware of node i's block at time t, and $|\mathbf{U}_i(t)|_\pi$ the combined block generation rate of $\mathbf{U}_i(t)$. We have $|\mathbf{U}_i(0)|_\pi = \pi - \pi_i$ and $|\mathbf{U}_i(t)|_\pi = 0$ when t exceeds the minimum time needed for i's block to reach all nodes. The probability that the rest of network does not generate a competing block by time t can be written as:

$$P_{NC,i}(t) = \prod_{s=\delta}^{t} \left(1 - |\mathbf{U}_i(s)|_\pi\right) \tag{2}$$

Since $(1 - |\mathbf{U}_i(s)|_M)_{s=\delta}^{t}$ is an increasing sequence bounded by $(0, 1]$, thus $(P_{NC,i}(t))_{t=\delta}^{\infty}$ is a convergent sequence.

Then by the law of total probability, the average blockchain fork rate of the whole network is obtained by weighing $(1 - \lim_{t\to\infty} P_{NC,i}(t))$ with $P(M_i), \forall i$:

$$\begin{aligned} FR &= \sum_i P(M_i)\big(1 - \lim_{t\to\infty} P_{NC,i}(t)\big) \\ &= \sum_i \frac{\pi_i}{\pi}\Big(1 - \prod_{s=\delta}^{\infty}\left(1 - |\mathbf{U}_i(s)|_\pi\right)\Big) \end{aligned} \tag{3}$$

When $\pi \ll 1$, N is large (e.g. π=1/600, $N\approx$10,000 in Bitcoin), mining power and network connectivity are evenly distributed, we have $\pi_i = \frac{\pi}{N}$, $|\mathbf{U}_i(s)|_M = \frac{\pi}{N}|\mathbf{U}_i(t)| = \frac{\pi}{N}|\overline{\mathbf{U}}(t)|, \forall i$. Further assuming $\delta \to 0$, then (3) reduces into the following form:

$$FR \approx 1 - \Big(1 - \pi\Big)^{\int_0^\infty \frac{1}{N}|\overline{\mathbf{U}}(t)|dt} \tag{4}$$

which is consistent with the result obtained by Decker et al. [10]. The approximation $(1 - ax) \approx (1 - x)^a$ for small x is used.

B. Mining Revenue and Relative Mining Gain

Define a discrete-time random process $\{B_i(h)\}_{h=1,2,...}$ in which $B_i(h) = 1$ if node i is the block generator at height h in the canonized blockchain; 0 otherwise. The mining revenue MR_i and relative mining gain RMG_i of node i in the long term are defined as follows:

$$MR_i = \lim_{H\to\infty} \frac{1}{H}\sum_{h=1}^{H} B_i(h) \tag{5}$$

$$RMG_i = \frac{MR_i - \pi_i/\pi}{\pi_i/\pi} \tag{6}$$

Next we propose an estimation method for MR_i via probabilistic analysis. Define another discrete-time random process $\{W_i(c)\}_{c=1,2,...}$, which is embedded right after each blockchain canonization event. Therefore there is no outstanding fork nor propagating block in the network when random variables $W_i(c)|_{c=1,2,...}$ are evaluated. We further define $h(c)$ as the blockchain height of the c^{th} canonization event and

$$W_i(c) = \begin{cases} 1 & \textit{if } B(h(c)+1) = i \\ 0 & \textit{otherwise} \end{cases} \tag{7}$$

Next we argue that the expectation of $W_i(c)$ at any epoch c, denoted $E[W_i]$, can be used to estimate MR_i in a conservative manner.

***Proposition** 1: $W_i(c)|_{c=1,2,...}$ are independent and identically distributed (i.i.d.) and their common expectation $E[W_i]$ satisfies the following relation with MR_i:*

$$E[W_i] \begin{cases} \leq MR_i & \textit{if } E[W_i] \geq \frac{\pi_i}{\pi} \\ > MR_i & \textit{otherwise} \end{cases} \tag{8}$$

In other words, $E[W_i] - \pi_i/\pi$ is a conservative estimate of the mining gain/loss of node i. Moreover, the gap between $E[W_i]$ and MR_i tightens as the overall fork rate FR decreases.

A proof sketch: Since $\{W_i(c)\}_{c=1,2,...}$ is embedded right after each blockchain canonization event when all previous forks are pruned and block propagation ceases, the competition for future blocks is oblivious of the block competitions in the past. And block generation at each node is a memoryless process. Therefore, $W_i(c)|_{c=1,2,...}$ are i.i.d, and we are able to compute their common expectation, denoted $E[W_i]$.

For an arbitrary canonization interval $c \to c+1$, we consider the blocks within it: those of height $h(c)+1, ..., h(c+1)$. First, $\frac{\pi_i}{\pi}$ evaluates the chance of i being the first to generate a block. $E[W_i] > \frac{\pi_i}{\pi}$ implies that i has a communication advantage over the network average which brings it positive mining gain. If $E[W_i] > \frac{\pi_i}{\pi}$ and i wins block $h(c)+1$, it will continue with a higher chance of winning the subsequent blocks from $h(c)+2$ to $h(c+1)$ because of its enhanced communication advantage during the fork race. Conversely, if $E[W_i] < \frac{\pi_i}{\pi}$ and i does not win block $h(c)+1$, the chance for i to win any block from height $h(c)+2$ to $h(c+1)$ further decreases because of its aggravated communication disadvantage. In contrast, $W_i(c)$ only considers the first block after canonization event c, and using $E[W_i]$ to estimate MR_i would assume i would have equal chance of winning any subsequent block from $h(c)+2$ to $h(c+1)$ as it won $h(c)+1$. Therefore, $E[W_i]$ tends to underestimate (or overestimate) MR_i if $E[W_i]>$(or $<$)$\frac{\pi_i}{\pi}$.

On the positive side, if the fork rate decreases, so is the interval $c \to c+1$, and so is the block count $h(c+1)-h(c)$ within the interval. That is, there will be fewer blocks in a fork incident for $E[W_i]$ to under-/overestimate MR_i, and thus the former can achieve higher estimation accuracy. ■

Next we calculate $E[W_i]$. By the law of total expectation:

$$E[W_i] = P(M_i)E[W_i|M_i] + \sum_{j\neq i} P(M_j)E[W_i|M_j] \quad (9)$$

We separated the summation because the two conditional events $W_i|M_i$ and $W_i|M_j$ occur under different condition.

$W_i|M_i$ consists of two subcases:

- *No-fork win:* No conflicting blocks are proposed by the rest of the network during the propagation of node i's block.
- *Fork win:* Conflicting blocks are proposed by the rest of the network during the propagation of node i's block, whereas node i's block still wins.

The probability of no-fork win equals to $P_{NC,i}(\infty)$, as was evaluated by (3). The probability of fork win is slightly more complicated. During the propagation of node i's block, the number of conflicting blocks generated by the rest of network at time slot $(t, t+\delta]$ conforms to a Bernoulli distribution with rate $|\mathbf{U}_i(t)|_M$. If node j happens to generate a competing block during $(t, t+\delta]$, node i's block will need to win the support of the majority computing power of the network before it encounters node j's block in a stalemate. We denote the chance of node i winning the fork under this condition by $\omega_{i\succ j}(t)$. Therefore:

$$\begin{aligned} &E[W_i|M_i] \\ &= E[W_i, \textit{No-fork win}|M_i] + E[W_i, \textit{Fork win}|M_i] \\ &= \lim_{t\to\infty} P_{NC,i}(t) + \sum_{t=\delta}^{\infty} P_{NC,i}(t) \sum_{j\in\mathbf{U}_i(t)} \pi_i \omega_{i\succ j}(t) \end{aligned} \quad (10)$$

Notably, in the derivation above we only considered two-tine forks for simplifying analysis; the likelihood of three or more-tine forks is negligible compared to that of two-tine forks.

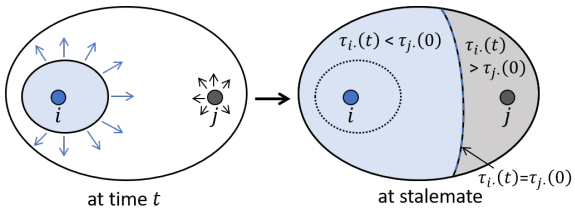

Fig. 3. Explanation of (12). Light blue (grey) area denotes portion of the network that advocates i's (j's) block. $\hat{\omega}_{i\succ j}(t)$ is evaluated by the total computing power covered by light blue area at stalemate.

In contrast, the conditional event $W_i|M_j$ in (9) can only happen via a fork race. That is, node i needs to generate a competing block during the propagation of node j's block, and eventually wins the fork. Similarly to (10), we have:

$$\begin{aligned} &E[W_i|M_j] \\ &= E[W_i, \textit{fork win}|M_j] \\ &= \sum_{t=\delta}^{\infty} P_{NC,j}(t)\pi_i \mathbb{1}_{\{i\in\mathbf{U}_j(t)\}}\big(1-\omega_{j\succ i}(t)\big) \end{aligned} \quad (11)$$

$\mathbb{1}_{\{i\in\mathbf{U}_j(t)\}}$ is an indicator function, returning 1 if the condition holds true; 0 otherwise. The winning chance of node i under this circumstance is $1-\omega_{j\succ i}(t)$.

Evaluating $\omega_{i\succ j}(t)$**.** $\omega_{i\succ j}(t)$ essentially measures the communication advantage of node i over node j when j generates a competing block. For i to win the fork race against j, it has to have the majority of the network advocate its block before the two competing blocks end up in a stalemate. Let the moment when i publishes its block be time 0. Define $\tau_{ik}(t)$ as the minimum time for i's block to propagate to node k starting from time t. Then $\omega_{i\succ j}(t)$ can be evaluated as follows:

$$\hat{\omega}_{i\succ j}(t) = \sum_k \pi_k \big(\mathbb{1}_{\{\tau_{ik}(t)<\tau_{jk}(0)\}} + \frac{1}{2}\mathbb{1}_{\{\tau_{ik}(t)=\tau_{jk}(0)\}}\big) \quad (12)$$

Figure 3 explains the calculation of $\hat{\omega}_{i\succ j}(t)$. As a result, we can finally obtain $E[W_i]$ by substituting (1), (10), (11), (12) into (9).

C. Security Analysis

We consider all nodes are honest-but-potentially-colluding. The fork rate *FR* provides an overall measure of how much mining power is wasted, while the 50%-attack threshold *AT50* measures the system's security in the worst case scenario that the colluding group consists of the highest mining revenue earners. Next we analyze how network connectivity impacts *FR* and *AT50*.

1) Lower overall network connectivity leads to higher FR and lower AT50, thus weaker consensus security: We assume the block generation rate π_i is fixed for any node i. First, lower overall network connectivity means it takes longer for any node to disseminate a new block across the network. This can be caused by a protocol change that lowers the minimum peer number requirement. As for the calculation in (2) (3), this leads to a higher $|\mathbf{U}_i(s)|_\pi$, a lower $\lim_{t\to\infty} P_{NC,i}(t), \forall i$, and thus a higher *FR*. Moreover, a lower $\lim_{t\to\infty} P_{NC,i}(t)$ means that more of MR_i comes from fork races and the distribution of mining revenue is deeper influenced by each node's communication capability. As a result, MR_i moves farther from $\frac{\pi_i}{\pi}$ and *AT50* moves lower.

Notably, for a certain network connectivity profile, higher block generation rate across all nodes (thus a higher π) would lead to a higher $|\mathbf{U}_i(s)|_\pi, \forall i$ and have the same impact of lower overall network connectivity.

2) Higher heterogeneity of network connectivity also leads to lower AT50: We still assume the block generation rate $\pi_i, \forall i$ is fixed. Higher heterogeneity of network connectivity means there is a greater divergence of communication capability among nodes. For instance, if node i resides in a highly-connected cluster in the network while node j resides in a sparsely-connected region, i will have a significant communication advantage over nodes in the sparse network region including j. As a result, i can disseminate a block to majority of the network much faster than j. $\hat{\omega}_{i\succ j}(t)$, as is evaluated by (12), will be close to 1 and $\hat{\omega}_{j\succ i}(t)$ will be much lower than 0.5. Therefore, i can harvest more mining revenue from fork races than j or other nodes in sparse network region. Consequently, $E[W_i]$ climbs higher above $\frac{\pi_i}{\pi}$ and $E[W_j]$ drops lower below $\frac{\pi_j}{\pi}$. This ultimately results in a more unequal *MR* distribution and hence lower *AT50*.

V. Incorporating Network Connectivity into Selfish Mining Analysis

In this section we evaluate the impact of network connectivity on selfish mining pool's communication capability and analyze its security implication under an expanding-consortium setting.

A. Selfish Mining Strategy

The core idea of selfish mining is to withhold newly generated blocks in a private chain, and release the private chain when the selfish mining pool sees the honest chain catch up close enough with the private chain. The detailed selfish mining strategy is illustrated in Figure 4, which we replicated from [11] and added with more description. Let α, β be the computing power share of the selfish mining pool and the honest nodes. Then $\alpha = \pi_{SM}/\pi$ and $\beta = 1-\alpha$, where π_{SM} is the selfish mining pool's block generation rate. The state number denotes the private chain's lead over the honest chain. State transition is triggered by any block generation event. Transitions from state 1 to $0'$ and 2 to 0 are accompanied by the selfish mining pool releasing the private chain. Any transition destined to state 0 marks a canonization

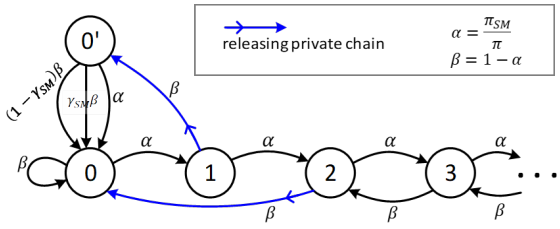

Fig. 4. Selfish mining strategy in [11]. State number denotes the private chain's lead over the honest chain. State transition is triggered by block generation.

event. γ_{SM} is defined as the long-term average fraction of honest computing power that will advocate the selfish mining pool's private chain when the pool and an honest miner release competing blocks simultaneously.

To incorporate network connectivity into the analysis, we model the selfish mining pool's network function as follows:

- Information exchanges within the selfish mining pool are instantaneous. The pool members are fully connected and synchronized. Any pool member who receives a new block from an honest node can make decision (changing state, switching chain, publishing the private chain) on behave of the entire pool.
- Once the selfish mining strategy determines to release the private chain, all pool members release the private chain to all peers simultaneously.
- Selfish mining pool members still relay blocks for honest miners, as long as the block does not trigger the pool to release its private chain. The reason is two-fold for the pool: to avoid suspicion of being a "blackhole" attacker, and to avoid network partitioning which would paralyze the blockchain system altogether.

Based on this model, we consider γ_{SM} the selfish mining pool's communication capability measure and evaluate it from the network connectivity profile.

B. Evaluating γ_{SM} Using Betweenness Centrality

Based on the assumption that nodes in the selfish mining pool are synchronous and can communicate with each other instantaneously, we treat these pool nodes as a fully-interconnected cluster and equivalently a super node denoted by SMPOOL, which preserves the pool members' all external communication links to the remaining honest nodes. We show that a betweenness centrality measure of SMPOOL within the network accurately evaluates γ_{SM}.

***Proposition** 2: γ_{SM} can be evaluated by the mining power weighted betweenness centrality measure of* SMPOOL*:*

$$\gamma_{SM} = \sum_{i\neq j\neq \text{SMPOOL}} \frac{\sigma(i,j|\text{SMPOOL})}{\sigma(i,j)} \cdot \frac{\pi_i}{\pi-\pi_{SM}} \cdot \frac{\pi_j}{\pi-\pi_{SM}-\pi_j} \tag{13}$$

wherein $\sigma(i,j)$ is the number of shortest paths between i and j, and $\sigma(i,j|\text{SMPOOL})$ is the number of such paths that pass through SMPOOL.

A proof sketch: Let i and j denote a pair of honest nodes, with i being the miner of a new block which triggers SMPOOL to release its private chain according to the selfish-mine strategy. For j to switch to SMPOOL's private chain instead of accepting i's new block, the highest block of SMPOOL's private chain must be propagated to j before i's new block. In the graph model, this is necessitated by SMPOOL residing on a shortest communication path between i and j. Therefore, $\frac{\sigma(i,j|\text{SMPOOL})}{\sigma(i,j)}$ gives the likelihood that SMPOOL delivers its private chain to j ahead of i's block. The weight $\frac{\pi_i}{\pi-\pi_{SM}} \cdot \frac{\pi_j}{\pi-\pi_{SM}-\pi_j}$ evaluates the pair $\langle i,j \rangle$'s mining power contribution to γ_{SM} among all pairs of honest nodes. As a result, the mining power weighted betweenness centrality measure of SMPOOL computes the average fraction of honest mining power that will advocate SMPOOL's block after SMPOOL releases its private chain, thus accurately evaluates γ_{SM}. ■

Equation (13) can be conveniently computed with the Brandes algorithm [17]. If there are M honest nodes and they have equal mining power, i.e. $\pi_i = \frac{\pi-\pi_{SM}}{M}, \forall i \neq SM$, then the weight $\frac{\pi_i}{\pi-\pi_{SM}} \cdot \frac{\pi_j}{\pi-\pi_{SM}-\pi_j}$ becomes $\frac{1}{M(M-1)}$ and (13) reduces to the standard normalized betweenness centrality measure.

With γ_{SM} obtained, the calculation of the selfish mining pool's mining revenue follows the procedure of [18]. Notably, the mining revenue of pool is proportional to γ_{SM}.

C. Security Analysis

We consider the selfish mining pool as an expanding consortium among the network of honest nodes. Under this setting, we discuss how network connectivity affects γ_{SM} and consensus security w.r.t. security thresholds *AT50* and *PRTH*.

1) Lower overall network connectivity leads to higher γ_{SM}, lower AT50, and lower PRTH, thus less secure against selfish mining: Lower overall network connectivity leads to reduced communication capability of both the selfish mining pool and honest nodes. However, since the selfish mining pool consists of originally honest nodes and preserve all their external communication links with the remaining honest nodes, communication capability reduction of an average honest node will be more significant than that of the selfish mining pool. Therefore, SMPOOL will be residing in the shortest communication paths of more honest pairs, yielding a higher γ_{SM} for every α value. Consequently this yields lower *AT50* and *PRTH*.

2) Compared to AT50, PRTH is more sensitive to network connectivity changes: According to [11], *PRTH* is reached much earlier than *AT50* for any γ_{SM}. Due to the gradualism of selfish mining pool's expansion, the rate of the selfish mining pool harvesting new external communication links initially increases, then gradually slows down as the pool takes in more nodes. Thus as α increases from 0 to 50%, γ_{SM} grows quickly at first then slower as it becomes closer to 1. Also the mining revenue of selfish mining pool is proportional to γ_{SM}. Therefore, a moderate reduction of overall network connectivity would lead to significant decrease in *PRTH*, but limited decrease in *AT50*. *PRTH* is also a more practical security measure than *AT50* in the sense that once the pool size hits *PRTH*, joining the pool will be financially attractive to the remaining honest nodes in the network.

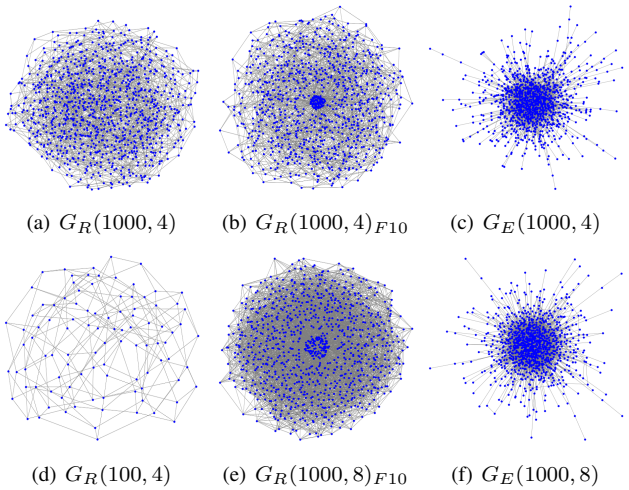

(a) $G_R(1000,4)$ (b) $G_R(1000,4)_{F10}$ (c) $G_E(1000,4)$

(d) $G_R(100,4)$ (e) $G_R(1000,8)_{F10}$ (f) $G_E(1000,8)$

Fig. 5. Visualization of six network graphs used in our experiments. Dot represents mining node, grey line segment represents communication link.

3) Selfish mining pool can take advantage of heterogeneity of network connectivity to achieve lower AT50 and PRTH: If the selfish mining pool is aware of the peer-to-peer network's topology, it can prioritize its expansion into well-connected regions of the network to maximize the growth of γ_{SM} and its mining revenue. As a result, heterogeneity of network connectivity can be exploited by selfish mining pool to achieve lower *AT50* and *PRTH*.

VI. Evaluation

We conducted simulation experiments to validate our model and security analysis. The simulation program was written in Python and follows a time-driven fashion and takes the following as input: graph representation of the network, block generation rates of all nodes, adversarial setting (honest or selfish mining), and simulation time (slots).

A. Setup

We use three types of network graph for evaluation with the following notations:

- $G_R(N,D)$: a regular graph with N nodes and degree D.
- $G_R(N,D)_{FX}$: a $G_R(N,D)$ but with the first $X\%$ of nodes being fully-interconnected.
- $G_E(N,D)$: a graph with N nodes and exponentially distributed node degrees with an average D.

The latter two graph types are designed to simulate different heterogeneous network connectivity profiles. The six network graphs used in our experiments are visualized in Figure 5.

To focus on evaluating the impact of network connectivity and provide a fair ground for comparing security thresholds, we assign all nodes the same block generation rate: $\pi_i = \frac{\pi}{N}, \forall i$ while using π as a variable to represent the aggregate block generation rate.

The follow metrics are used to evaluate our model and analysis: *FR*, *AT50* and *PRTH* for security metrics, rooted mean

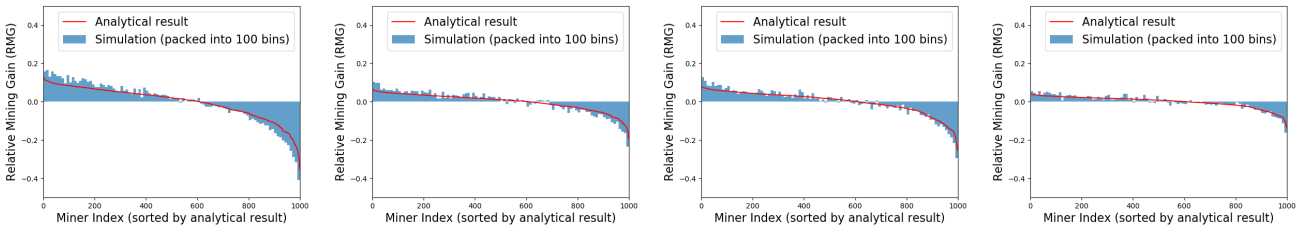

(a) $G_E(1000, 4), \pi = 0.1$ (b) $G_E(1000, 4), \pi = 0.05$ (c) $G_E(1000, 8), \pi = 0.1$ (d) $G_E(1000, 8), \pi = 0.05$

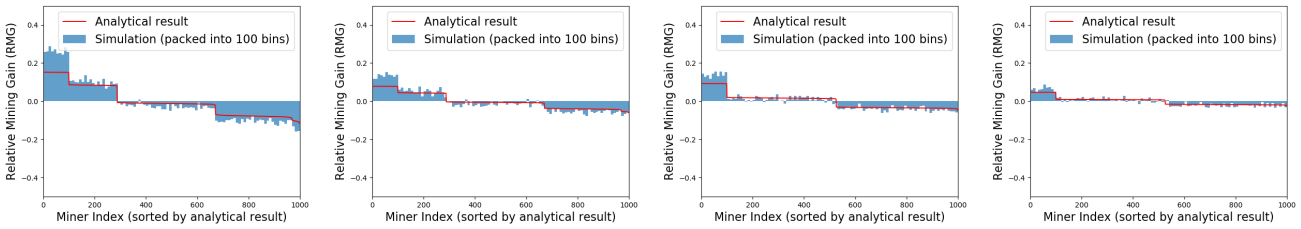

(e) $G_R(1000, 4)_{F10}, \pi = 0.1$ (f) $G_R(1000, 4)_{F10}, \pi = 0.05$ (g) $G_R(1000, 8)_{F10}, \pi = 0.1$ (h) $G_R(1000, 8)_{F10}, \pi = 0.05$

Fig. 6. Relative mining gain (*RMG*) results of eight experiments.

TABLE II
HONEST MINING EXPERIMENT RESULT CORRESPONDING TO FIGURE 6

	Configuration		**Metrics**				
	Graph(N, D)	π	FR-SIM	FR-ANA	AT50-SIM	AT50-ANA	RMSE
a	$G_E(1000, 4)$	0.1	0.3148	0.3136	459/1000	470/1000	0.0325
b	$G_E(1000, 4)$	0.05	0.1773	0.1670	478/1000	484/1000	0.0205
c	$G_E(1000, 8)$	0.1	0.2409	0.2248	475/1000	479/1000	0.0187
d	$G_E(1000, 8)$	0.05	0.1315	0.1159	487/1000	489/1000	0.0123
e	$G_R(1000, 4)_{F10}$	0.1	0.3117	0.3099	457/1000	470/1000	0.0423
f	$G_R(1000, 4)_{F10}$	0.05	0.1758	0.1649	479/1000	485/1000	0.0232
g	$G_R(1000, 8)_{F10}$	0.1	0.2309	0.2124	480/1000	484/1000	0.0195
h	$G_R(1000, 8)_{F10}$	0.05	0.1250	0.1090	490/1000	491/1000	0.0113

square error (*RMSE*) between analytical *RMG* distribution and simulated *RMG* distribution for the model accuracy metric.

B. Honest Mining Experiment

We performed eight experiments on four network graphs with different settings. Each experiment was run for 10 million timeslots. The configuration and evaluation results are shown in Table II and the *RMG* histograms are shown in Figure 6. We made the following observations:

1) The analytical RMG result conservatively estimates the simulation result. The accuracy improves when π decreases or D increases. The obvious gap between analytical result and simulation in Figure 6(c) demonstrates the fully-interconnected top-10% have a significant higher block winning chance than that expected by $E[W_i]$. Nonetheless, as is shown in Table II, for either graph type when π decreases from 0.1 to 0.05 or D increases from 4 to 8 the fork rate decreases and so does *RMSE*. This validates Proposition 1 that the estimation accuracy improves when fork rate drops.

2) FR decreases and AT50 increases when π decreases or D increases. This validates the security analysis in IV-C in that higher overall network connectivity or lower block generation rate leads to stronger consensus security in the presence of potentially colluding nodes.

C. Selfish Mining Experiment

We switched the adversary setting from honest to selfish mining and performed three experiments. Each experiment targeted a certain network graph and consisted of seven simulations, each took an α value from $\{2, 5, 10, 20, 30, 40, 45\}\%$ and ran for 10 million timeslots. Figure 7(a) shows the analytical result of γ_{SM}. The configuration and evaluation results are shown in Table III and Figure 7(b). For graph $G_E[1000, 4]$ we configured the selfish mining pool to expand from the highest-degree node to lower-degree nodes in a descending order. This was purposed as an example of selfish mining pool's expansion strategy. *RMSE* here measures the averaged estimation accuracy of the analytical *RMG* over all α values.

To estimate the thresholds *AT50* and *PRTH*, for each of the three experiments we fitted the simulated *RMG* points using degree-7 polynomials and obtained *AT50* and *PRTH* using the fitted curve. The following observations are made:

1) The analytical result matches simulation. The close match between analytical *RMG* and simulation in Figure 7(b) validates our betweenness centrality-based calculation of γ_{SM}.

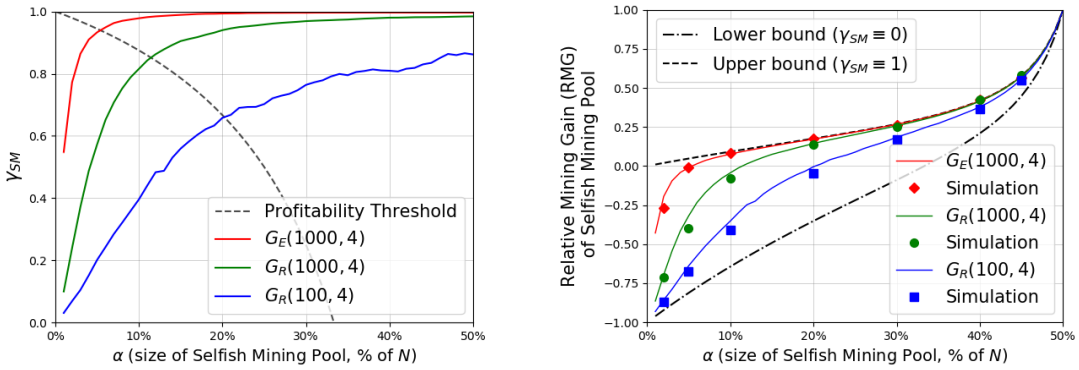

(a) Analytical γ_{SM}.

(b) Analytical RMG and simulation.

Fig. 7. Comparison of simulation and analytical result for selfish mining.

TABLE III
SELFISH MINING EXPERIMENT RESULT CORRESPONDING TO FIGURE 7

	Configuration		Metrics				
	Graph (N, D)	π	PRTH-SIM	PRTH-ANA	AT50-SIM	AT50-ANA	RMSE
1	$G_E(1000, 4)$	0.01	52/1000	55/1000	369/1000	369/1000	0.0290
2	$G_R(1000, 4)$	0.01	122/1000	114/1000	368/1000	370/1000	0.0338
3	$G_R(100, 4)$	0.01	22/100	21/100	38/100	38/100	0.0311

2) When N decreases from 1000 to 100, PRTH changes more dramatically than AT50. This validates our security analysis that *PRTH* is more sensitive to network connectivity changes. Particularly, the analytical curves of γ_{SM} in Figure 7(a) demonstrates that as the selfish mining pool size expands from $\alpha = 0$ to $\alpha = 50\%$, γ_{SM} grows quickly at first then gradually slows down when it comes closer to 1.

3) As is shown in Figure 7(a), γ_{SM} rapidly crosses the profitability threshold and grows close to 1 when the selfish mining pool expands in $G_E(1000, 4)$. This demonstrates the feasibility for selfish mining pool to choose a particular expansion strategy to exploit heterogeneity of network connectivity for faster revenue growth.

VII. Discussion and Future Work

On Potential Model Deficiency In our model we only consider two-tine forks. That is, at most two blocks of the same height could be propagating in the network concurrently. Though the possibility of three-tine forks or more is significantly lower than two-tine forks, it could still affect the long-term mining revenue distribution, especially when the network is large and block propagation delays are high. To tackle this issue we would need a more delicate model that takes into account the subtleties of the progression of a fork race. We leave it to future work.

On Model Practicality In practical networks, it can be challenging to monitor the block propagation progress (i.e. $U_i(t)$). To overcome this difficulty, a block propagation progress-agnostic model is needed to estimate the communication capability and forecast the revenue of a node via congregate network statistics. Furthermore, practical PoW blockchain networks tend to be structurally volatile and may conform to a scale-free growth pattern [3]. To take that into account, it is possible to model structural changes of the network with a certain random process and evaluate its impact on information propagation and consensus security. The impact of the selfish mining pool's internal communication routine and structural dynamicity is also a potential issue to explore.

VIII. Conclusion

We presented a modeling study on the impact of network connectivity on consensus security of PoW blockchain under two adversarial scenarios, namely honest-but-potentially-colluding and selfish mining. For the first scenario, we demonstrated the communication advantage of a node over its competitors in a fork race and provided a method to estimate its long-term mining revenue and relative mining gain. For the second scenario, we introduced a practical model for the selfish mining pool's network functions and showed that the communication capability of selfish mining pool, γ_{SM}, can be accurately evaluated by the mining power-weighted betweenness centrality measure. For both scenarios, we showed that low network connectivity and excessive heterogeneity of network connectivity lead to poor consensus security. Our analysis can serve as a paradigm for associating consensus security with network connectivity. In future work we will incorporate more realistic network settings into our model and extend the analysis to other blockchain consensus protocols.

Acknowledgment

This work was supported in part by US National Science Foundation under grants CNS-1916902 and CNS-1916926.

REFERENCES

[1] S. Nakamoto *et al.*, "Bitcoin: A peer-to-peer electronic cash system," 2008.

[2] J. Garay, A. Kiayias, and N. Leonardos, "The bitcoin backbone protocol: Analysis and applications," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pp. 281–310, Springer, 2015.

[3] A. Baumann, B. Fabian, and M. Lischke, "Exploring the bitcoin network.," in *WEBIST (1)*, pp. 369–374, 2014.

[4] A. Miller, J. Litton, A. Pachulski, N. Gupta, D. Levin, N. Spring, and B. Bhattacharjee, "Discovering bitcoin's public topology and influential nodes," *et al*, 2015.

[5] T. Neudecker, P. Andelfinger, and H. Hartenstein, "Timing analysis for inferring the topology of the bitcoin peer-to-peer network," in *2016 Intl IEEE Conferences on Ubiquitous Intelligence & Computing, Advanced and Trusted Computing, Scalable Computing and Communications, Cloud and Big Data Computing, Internet of People, and Smart World Congress (UIC/ATC/ScalCom/CBDCom/IoP/SmartWorld)*, pp. 358–367, IEEE, 2016.

[6] A. E. Gencer, S. Basu, I. Eyal, R. Van Renesse, and E. G. Sirer, "Decentralization in bitcoin and ethereum networks," *arXiv preprint arXiv:1801.03998*, 2018.

[7] Y. Sompolinsky and A. Zohar, "Accelerating bitcoin's transaction processing," *Fast Money Grows on Trees, Not Chains*, 2013.

[8] Y. Sompolinsky and A. Zohar, "Secure high-rate transaction processing in bitcoin," in *International Conference on Financial Cryptography and Data Security*, pp. 507–527, Springer, 2015.

[9] 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 *International Conference on Financial Cryptography and Data Security*, pp. 106–125, Springer, 2016.

[10] C. Decker and R. Wattenhofer, "Information propagation in the bitcoin network," in *Peer-to-Peer Computing (P2P), 2013 IEEE Thirteenth International Conference on*, pp. 1–10, IEEE, 2013.

[11] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," *Communications of the ACM*, vol. 61, no. 7, pp. 95–102, 2018.

[12] A. Sapirshtein, Y. Sompolinsky, and A. Zohar, "Optimal selfish mining strategies in bitcoin," in *International Conference on Financial Cryptography and Data Security*, pp. 515–532, Springer, 2016.

[13] 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 *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 3–16, ACM, 2016.

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

[15] S. Bano, A. Sonnino, M. Al-Bassam, S. Azouvi, P. McCorry, S. Meiklejohn, and G. Danezis, "Consensus in the age of blockchains," *arXiv preprint arXiv:1711.03936*, 2017.

[16] Y. Xiao, N. Zhang, W. Lou, and Y. T. Hou, "A survey of distributed consensus protocols for blockchain networks," *arXiv preprint arXiv:1904.04098*, 2019.

[17] U. Brandes, "A faster algorithm for betweenness centrality," *Journal of mathematical sociology*, vol. 25, no. 2, pp. 163–177, 2001.

[18] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," in *International conference on financial cryptography and data security*, pp. 436–454, Springer, 2014.