Towards a Model for Comprehending and Reasoning about PoW-based Blockchain Network Sustainability

Sotirios Liaskos
School of Information Technology
York University
Toronto, Canada
liaskos@yorku.ca

Bo Wang
Department of Electrical Engineering
and Computer Science
York University
Toronto, Canada
bowang@eecs.yorku.ca

ABSTRACT

Blockchain networks have been claimed to have the potential of fundamentally changing the way humans perform economic transactions with each other. In such networks, trust-enabling agents and activities, that were traditionally arranged in a centralized fashion, are replaced by a network of nodes which collectively yet independently witness and establish the non-repudiability of transactions. Most often, a proof-of-work (PoW) requirement ensures that participants invest resources for joining the network, incentivizing conformance to the network rules, while making it highly infeasible for malicious agents to construct an alternative version of the transaction history. While research on security and efficiency aspects of blockchain networks is already being conducted, there is still work to be done to understand how different external and internal conditions guarantee or threaten their sustainability, i.e., their continuous operation. Focusing on public PoW-based blockchain platforms, in this paper we sketch an abstract model that is aimed at supporting comprehension and qualitative reasoning about the factors that affect sustainability of a blockchain network.

CCS CONCEPTS

• **Applied computing** → **Digital cash**; **Electronic funds transfer**; • **Computer systems organization** → Peer-to-peer architectures;

KEYWORDS

blockchain networks, bitcoin, cryptocurrencies

ACM Reference Format:
Sotirios Liaskos and Bo Wang. 2018. Towards a Model for Comprehending and Reasoning about PoW-based Blockchain Network Sustainability. In *Proceedings of ACM SAC Conference (SAC'18).* ACM, New York, NY, USA, Article 4, 5 pages. https://doi.org/https://doi.org/10.1145/3167132.3167175

1 INTRODUCTION

Since their introduction a few years ago, blockchain networks have enjoyed considerable attention by the technical and business community, the media and the society at large. Such networks aim at offering a non-repudiable way of maintaining transactions between social and economic agents. In blockchain networks, the party that witnesses and establishes the validity of such transactions is not a unique externally validated business or legal entity but rather a community of anonymous antagonizing computational agents [1]. This arrangement is believed to have the potential to bring about numerous and profound implications in modern societies, from increasing transaction security and dependability thanks to the transition from one to multiple points-of-failure to bringing about inclusiveness, transparency and democratization in social and economic affairs through disallowing concentration of validation power to a few large organizational and legal entities [9].

The sustaining successful operation of Bitcoin [8], the cryptocurrency that introduced the technology, has allowed the community to maintain considerable confidence on the soundness of the underlying technical design of blockchain platforms. Nevertheless, disproportionately little independent research has been conducted on conceptualizing and analyzing blockchain networks as complex systems and exploring what conditions, internal and external, affect their behavior and ultimately their long-term sustainability and utility.

In this paper, we attempt a simple model for assisting comprehension of and qualitative reasoning about the sustainability of public proof-of-work (PoW-based) blockchain networks. These are blockchain networks that feature a volatile and large set of participants and use search-based computation as the proof-of-work method for excluding disruptive and malicious participants. Bitcoin is the archetypal PoW-based public blockchain. The proposed model focuses on the parameters that affect the decision of an individual node to participate in transaction validation. Extrapolating such a local decision model to the entire network our simple model aims at understanding how various configurable and observed network parameters interact with each other and how such understanding can help us re-configure the network's protocol to support sustainability in response to external or internal disruptions. We particularly show how sustainability may be affected by externally caused drops in its token price and what parameter reconfiguration can potentially prevent node disincentivization and departure.

We organize our presentation as follows: Section 2 offers an introduction to blockchain networks, Section 3 presents the model, Section 4 demonstrates uses of the model and Sections 5 and 6 offer concluding remarks and pointers to related work.

SAC'18, April 9-13, 2018, Pau, France

ACM ISBN 978-1-4503-5191-1/18/04…$15.00
https://doi.org/https://doi.org/10.1145/3167132.3167175

2 PUBLIC POW-BASED BLOCKCHAINS

Blockchain networks consist of a set of interconnected nodes, each maintaining a copy of the history of all transactions that have taken place within the purview of the network (the *blockchain*). Transactions arrive at nodes and are immediately propagated through the network. The transactions represent exchange of *tokens* which are units of value intrinsic to the network, but with an external price as well, measured in traditional currencies.

Nodes organize the transactions they receive into *blocks*, validate them, append them to the blockchain and propagate this blockchain augmentation to all other nodes. However, the validation rules are such that, in addition to checking the correctness of the included transactions, nodes need to carry out a considerable amount of computation as *proof-of-work* that the new block comes from a legitimate non-disruptive node. By demanding that nodes commit to large-scale computation and dedicate the appropriate energy and infrastructure, the network ensures that only nodes that are invested in the sustainability of the network participate. At the same time, the development and adoption of an alternative version of the transaction history – i.e. an alternative competing blockchain by, e.g., a malicious node or group thereof, would require computational power comparable to that of the entire network, a goal understood to be highly impractical to attain.

In Bitcoin and other public blockchain networks, the proof-of-work computation is materialized as a brute-force search within a large search space to find an element that belongs to a specific subset. The size of the subset relative to the size of the search space is an indication of the *difficulty* of the proof-of-work aspect of the network. The proof-of-work computation is translated into energy and infrastructure investment, which is understood to be the cost the node has to pay to participate in the network.

Nodes are, nevertheless, motivated to participate thanks to a reward, in form of network tokens, that is given to them for successfully validating a block before other nodes do. Transaction fees specified in individual transactions included in the block are also added to the block reward. Given the associated costs and rewards a foundational sustainability condition for PoW public blockchain networks can be drawn: within a given period of time, for every node, the expected external value of the reward must be at least as much as the cost dedicated to validation computation with that time.

Failure to meet the sustainability condition for an anticipated period of time must be assumed to lead any rational node to the decision to opt-out of the network for that time. But how do factors such as PoW difficulty, token price, computation cost, power distribution and network size conditions affect this cost-benefit equation? Importantly how do such opt-out opt-in decisions happening en masse affect the network characteristics? The simple model we develop below aims at allowing comprehension of the interplay between these factors and prepare the ground for more refined predictive modeling.

3 TOWARDS A SUSTAINABILITY MODEL

PoW Computation. The PoW computation of a public blockchain is defined within a search space S of size $|S|$. To attain the computational goals of PoW each node must search this space to find elements that belong to a specific subset $S_s \subseteq S$. The size of the subset in comparison to the overall search space describes the *difficulty*[1] d of the PoW $|S|/|S_s|$. To allow for predictability of search time, the search problem is such that only brute force search of the space is meaningful, entailing a procedure of linearly picking values from S and testing if they are included in S_s. We refer to each act of identifying and testing such value as a *trial* (in Bitcoin jargon: a hash).

Nodes, Cost and Revenue. A blockchain network consists of a set of validating nodes $n \in N$ capable of performing PoW computations (in Bitcoin: miners). Based on their hardware capability, the nodes are able to perform a number of trials per second at a rate which constitutes the power p_n (in Bitcoin: hash-power) of the node – measured as trials/sec. The node is also characterized by a cost factor c_n measured as the external cost of each trial (e.g., USD/trial). In pragmatic terms the cost includes computational and cooling energy as well as infrastructure leasing or depreciation. It follows that in a period of time t a node utilizing power p_n with cost factor c_n assumes total cost:

$$C_n = t \cdot p_n \cdot c_n \tag{1}$$

Nevertheless, the network protocol allows for a reward of r network tokens for successful validation of a block. Each token has an extrinsic value x measured in a traditional currency, e.g., USD. Thus, if successful validation within time t happens with probability pr_n, the node will accrue a total expected reward of:

$$R_n = pr_n \cdot r \cdot x \tag{2}$$

Competitive Search. Every time a node is engaged in validating a block, it competes to outperform all other nodes in the network working on similar blocks. Once a winner is declared a new competition (*game*) starts with a new block. The blocks are similar in that the sets of transactions they include are most likely to have a non-empty intersection. As each transaction can be validated only once, arrival of news that a similar block has been validated by a competing node implies failure for the current game. The set of competing nodes can be modeled as a substantially more powerful computer – one that utilizes massive parallelism.

When a node has performed a number of trials tr_n, there is a certain probability that the node has found one or more solutions that satisfy the conditions for being included in S_s, before the rest of the network does. The probability pr_n by which this happens is related to the power ratio between the node and the network. In particular, assume that the node and the network have been competing with each other for time Δt, when a winner is found. During that time the node has performed tr_n trials and the network tr_w trials. For our purposes, we can approximate the probability that the wining trial is one that the node performed by $pr_n = tr_n/(tr_n + tr_w) = p_n/(p_n + p_w)$ and thus:

$$pr_n = p_n/w \tag{3}$$

where $w = p_w + p_n$ is the power of the entire network.

Competition Period. The expected duration of a competition between the node and the network manifests itself as the average time between two consecutive block validations, the *block time* t_b. This time can be directly observed by comparing timestamps in the

[1] A concept similar but not mathematically identical to Bitcoin's difficulty.

blockchain. Moreover, the blockchain governance may be such that the difficulty level is adapted so that the interval between block validations stays, with some error, at a certain level – as in Bitcoin.

The observed time between blocks can be used to estimate the power of the network. The average trials within S to successfully draw an item from $S_s \subseteq S$, $|S_s| = |S|/d$ is $|S|/|S_s| = d$ (trials). If the network can achieve it within the block time t_b, then an estimation of the total network power w is:

$$w = d/t_b \tag{4}$$

Given the above, the expected duration of a game is d/w and the probability of node n to win it is, as we saw, p_n/w for any estimation w of network power exerted during the game. Thus, node n's expected reward for each game becomes $R_n = (p_n/w) \cdot r \cdot x$, while the cost is $C_n = (d/w) \cdot p_n \cdot c_n$.

Sustainability. Given the above formulations we can set up a basic node-level sustainability model as follows:

$$\boxed{r \cdot x \geq c_n \cdot d} \tag{5}$$

Thus, for an arbitrary game given the cost and power characteristics of the node, expressed in c_n, as well as block reward r, token exchange rate x and difficulty level d, a rational node will play the game if the expected reward is at least as much as the expected cost, i.e., $R_n \geq C_n$. This simplifies into Eq. 5.

This simple model can now be used for various types of qualitative reasoning and simulation for exploring the interplay between internal and external parameters of the network, as we sketch in the next section.

4 APPLICATIONS

4.1 Node-Level Decisions

Consider a node fully utilizing validation equipment units[2] with power consumption rating at approx. 1323 W and promised power $p_n = 13.5$T trials/sec. We further assume access to electricity prices of \$0.1 per kWh ($\$2.78 \times 10^{-8}$ per Joule). Hence, the device has a cost factor of $c_n = 2.72 \times 10^{-18}$ (\$/trial). In this particular case, calculation of the total cost for t seconds of search per unit is easily done by $C_n = t \cdot 1323 \cdot 0.1/(3600 * 1000) = 3.68 \times 10^{-5} t$. Leasing and other infrastructure costs can be added – omitted here for simplicity.

The network exhibits a current configuration[3] which includes reward level $r = 12.5$, the search space size $|S| = 1.16 \times 10^{77}$ and current difficulty (according to our definition) level $d = 4.74 \times 10^{21}$. We further assume an average time between block validations of $t_b = 9.17$ mins. Given these the total power of the network can be calculated as $p_w = (d/t_b) - p_n = 8.62 \times 10^{18}$ trials/sec. Finally the current extrinsic price of the network token is \$3650.

The node needs to know whether it is worth turning on the validation unit under the given network conditions and token price. This can be done by solving Eq. 5 for x. The value we get, approx. \$1032 in our case, represents the token price below which turning on the unit is not a profitable decision. Thus, at a current electricity price of \$0.1, turning the unit on seems to be a profitable (in the long term) decision. Of course for real revenue to materialize in practical time (e.g., before conditions change unfavorably) many such units might need to work in full availability.

[2]Modeled after the Antminer S9 16nm ASIC Bitcoin Miner.

[3]All numbers are adapted approximations of Bitcoin parameters and observed characteristics as of Sept. 18th, 2017

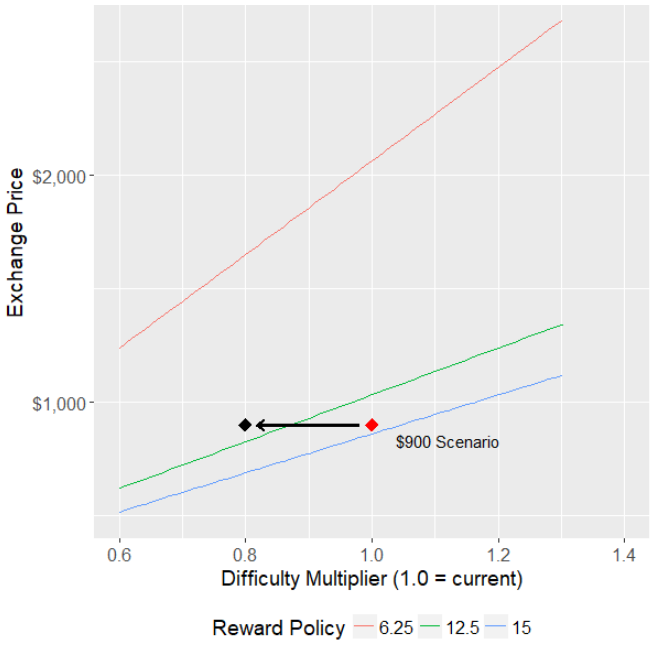

Figure 1: Price versus Network Power and Reward

4.2 Network Behavior

Perhaps more interesting than node-level reasoning is analysis at the network configuration level. Two main parameters are available for adjustment there: validation reward and difficulty level. In Bitcoin, for instance, while the latter is adjusted to maintain a relatively constant time between block validations, the former (reward) is adjusted in a rigid pre-programmed way (halved after a period of time). Reward and difficulty however both affect node-level decision.

Considering, again, the node we introduced earlier, increase in difficulty and/or reduction of reward and/or –as we saw– reduction of token exchange price may all contribute towards their decision to interrupt validation. Changes in difficulty are, in turn, caused by changes in the overall network power. In Figure 1, the combinations of exchange price and difficulty levels at which the node is indifferent as to whether they should participate are represented using three lines. Each line represents a different reward policy. Various difficulty levels are represented as multiplication factors against the current level (1.0). As we can see, if the token price became \$900, due to, e.g., an external economic event, our node would enter an unsustainable region.

Entering the unsustainability region does not affect only one node, but possibly a large proportion of the nodes in the network that operate under a similar cost factor. In our \$900 token price example, if, say, 20% of the network power consisted of nodes of similar characteristics, we would have a 20% network power drop.

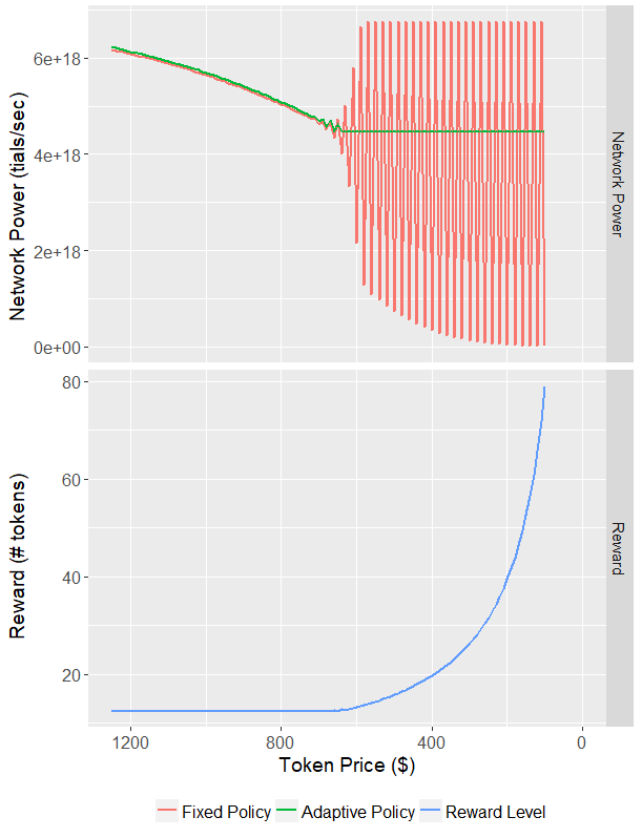

Figure 2: Price Drop Behavior

To investigate network behavior under these circumstances we can run a simple simulation. Let us consider a network consisting of 500,000 validation units with a normally distributed electrical power consumption (M = 1323 Watts, SD = 100 Watts), normally distributed power ($M = 1.35 \times 10^{13}$ trials/sec, $SD = 3 \times 10^{12}$ trials/sec), and access to, again, normally distributed electricity price (M = 0.1 \$/kWh, SD = 0.03 \$/kWh). The total amount of (hash-)power of the simulated network approaches the one we discussed earlier. For our example, we simulate a sequence of games in each of which the price of the token drops by \$10 from an initial value of \$1, 500 to a final value of \$100. Reward is fixed at $r = 12.5$. However, difficulty is adjusted in every step to keep block time t_b stable. Specifically, if total network power w can be estimated by dividing difficulty d by block time t_b (cf. Eq. 4), working reversely and assuming a value for w, the next difficulty level to sustain block time t_b assuming the same w can be set as $d \leftarrow w \cdot t_b$.

The result can be seen in the upper part of Figure 2, focusing on the red line labeled as "Fixed". At some price region around \$600, the sustainability frontier is crossed and many nodes who find the price unsustainable withdraw from the next round, leading to a drop in the network power, observed as an increase in block time. This, in turn, results in a decrease in difficulty for the next game as described above. Nevertheless, the decreased difficulty now lowers the break-even point and nodes that opted out in the previous round now decide to join again. The resulting increased power will recreate the original circumstances leading the system to the observed oscillation. Although oscillation may not be the necessary result – e.g., in cases where difficulty adaptation cannot catch-up with a rapid price drop – drop in the network power is rather inevitable if token price keeps decreasing.

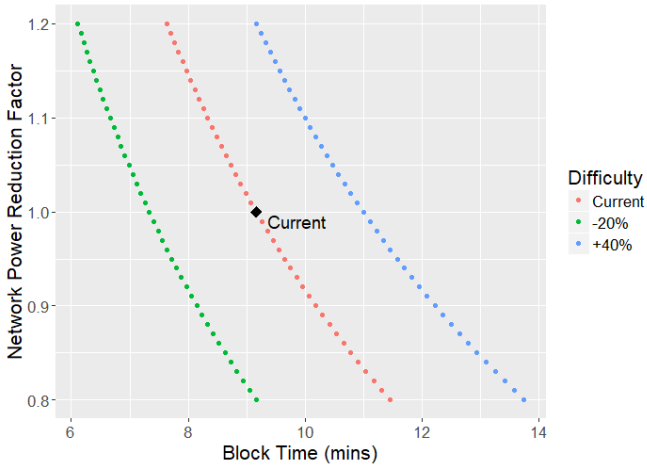

Figure 3: Power Fluctuations vis-a-vis Difficulty

Note, that, if difficulty is not adjusted immediately – e.g., in Bitcoin this only happens every 2016 blocks – our model suggests that a power drop will only express itself as prolonged block time. This can be seen in Figure 3. For as long as difficulty is fixed, changes in network power are expected to move along the corresponding line. As we saw, the hypothesis in the model is that profitability is not affected by smaller network power: the benefit from slightly easier games is balanced by the fact that games take now also a bit longer.

4.3 Reward Adaptation

Observe in Figure 1 that increasing the reward is another way to re-incentivize the nodes that opted out due to the price drop. In Figure 2 the green line labeled "Adaptive" represents the same price drop simulation we ran above, but with an adaptive reward policy. Specifically, to calculate the reward for the next game, we first calculate the break-even rewards of each node, i.e. the level of reward that would make the node indifferent with regards to whether it should participate in the next game. As we saw (Eq. 5), this depends on the difficulty level, the price and the cost factor. We calculate the 60% quantile of this set of break-even rewards, i.e. the reward level r_a below which 40% of the nodes will opt-out in the next round. Given a target reward r_t, in our case $r_t = 12.5$, in each round we set the reward to be $r \leftarrow max(r_t, r_a)$. Thus, when a substantial number of nodes is suspected to opt out of the next game, the reward increases to prevent this from happening. Otherwise, the reward stays in the target level to avoid unnecessary inflation.

In the figure, fixed and adaptive scenarios have the same behavior as long as more than 60% of the total nodes are motivated to participate. When this threshold is crossed due to the continuing

price decrease, the adaptation mechanism kicks in by increasing the reward level just as much as to keep no less than 60% of the total nodes motivated to continue participating. This results in stabilization of network power to that level. The lower part of the figure shows how the reward level increases as a result of the adaptation policy.

5 DISCUSSION AND CHALLENGES

The above are simple demonstrations of the kind of reasoning the proposed model allows with regards to the sustainability of public PoW-based blockchain networks. The model is deliberately abstract and not intended to be used for e.g., predictions. Rather, it aims at assisting comprehension and education and forming a basis for developing more refined quantitative predictive models in the future.

A particular point of interest is the choice of approximation of node success probability (Eq. 3) which allows us reason about node decision independent of certain parameters such as total network power (cf. Figure 3). In refining the model, analytical and empirical work will need to be performed for investigating formulations of the specific probability that are reasonably accurate for performing useful reasoning.

Moreover, we believe a case can be made for investigating reward and difficulty adaptation as a means for increasing the network's resilience to disruptive events. Nevertheless, there are many aspects of such adaptation techniques that are yet to be studied. The approach discussed above is based on the assumption that the size as well as power and cost factor distribution of the available set of nodes can somehow be estimated with relative accuracy, and that consensus can also be reached with regards to the amount by which the reward needs to be adjusted. Further, aspects such as the rate of the price change, its relationship to the overall token supply as well as the frequency of the difficulty and reward adaptation events need to be studied as parameters that affect adaptation properties. Moreover, a way to express adaptation objectives might need to be developed in terms of both aspired levels in parameters such as power, block time and total token supply and desired stability of those parameters.

Finally, pragmatic aspects need to be incorporated in such adaptation models. For instance, unless switching from validating one currency to another can be assumed [2], turning on and off validation rigs comes with a cost and a delay that needs to be taken into account in modeling node behavior. At the same time, node rationality needs to be assumed to be not only bounded by various levels of approximation of their cost factor but also dependent on strategic objectives, such as, for example, the possibility that a node thoroughly invested in the sustainability of the network will be ready to sustain losses for a period of time. Such factors are largely unobservable to an adaptation mechanism.

6 RELATED AND FUTURE WORK

Several models for modeling public blockchains have been attempted. E. Teo [10], for example, develops and simulates a model inspired by network economics, attempting to address the sustainability question by also incorporating concepts of risk level and mining pools – the extrinsic price of the token is not included in the analysis. Focusing on Bitcoin, Kroll et al. [7] also take a game-theoretic approach to sketch a sustainability equation of similar abstraction level like ours, offering also a discussion on central security threats for Bitcoin and how they affect incentives. N. Houy [6], focusing again on Bitcoin, discusses the interplay between block size, transaction fee constraints and mining rewards, through, again, taking a game-theoretic approach. Carlsten et al. [2] observe that in the absence of a block reward policy, where validation payment comes only in the form of transaction fees – which is also Bitcoin's planned long term state – there might be periods of time during which nodes loose incentive to participate due to lack of sufficient transaction fees to form profitable blocks, threatening, in other words, sustainability. Both in this context and elsewhere (e.g., [3, 5]) selfish mining behaviors become a central aspect of the analysis. Finally, a formal treatment of difficulty adjustment has also been offered by Garay et al. [4].

This body of work will inform candidate refinements and extensions of our model. The immediate item in our future research agenda is, nevertheless, the empirical validation of the proposed model and an examination of the strength of its assumptions and approximations. We anticipate that extended simulations combined with investigation of past data from real world crypto-currencies will offer us good starting points to attain such validity evidence. In addition, in further refining and validating our model, we plan to depart from Bitcoin-specific design concepts, which we adopted here for demonstration purposes, and make our model as generic as possible so that it is useful for analyzing a wider range of present or future blockchain network designs.

REFERENCES

[1] A. Narayanan, J. Bonneau, E. Felten, A. Miller, and S. Goldfeder. *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction.* Princeton University Press, 2016.

[2] M. Carlsten, H. Kalodner, S. M. Weinberg, and A. Narayanan. On the Instability of Bitcoin Without the Block Reward. In *Proc. of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS'16)*, pages 154–167, 2016.

[3] 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. of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS'16)*, pages 3–16, 2016.

[4] J. A. Garay, A. Kiayias, N. Leonardos. The Bitcoin Backbone Protocol: Analysis and Applications. In *Proc. of the 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques EUROCRYPT (2)*, pages 281–310, 2015.

[5] A. Kiayias, E. Koutsoupias, M. Kyropoulou, Y. Tselekounis. Blockchain Mining Games. In *Proc. of the 2016 ACM Conference on Economics and Computation (EC'16)*, pages 356–382, 2016

[6] N. Houy. The Bitcoin Mining Game. *Ledger*, 1:53–68, 2016.

[7] J. A. Kroll, I. C. Davey, and E. W. Felten. The Economics of Bitcoin Mining, or Bitcoin in the Presence of Adversaries. In *Proc. of the the 12th Workshop on the Economics of Information Security (WEIS 2013)*, pages 1–21, 2013.

[8] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. Technical report, www.bitcoin.org, 2008.

[9] D. Tapscott and A. Tapscott. *Blockchain Revolution: How the Technology Behind Bitcoin is Changing Money, Business, and the World.* Portfolio, 2016.

[10] E. G. S. Teo. Emergence, Growth, and Sustainability of Bitcoin: The Network Economics Perspective. In *Handbook of Digital Currency: Bitcoin, Innovation, Financial Instruments, and Big Data*, pages 191–200. Elsevier Inc., 2015.