



King's Research Portal

DOI:

[10.1109/TEM.2020.2981286](https://doi.org/10.1109/TEM.2020.2981286)

Document Version

Publisher's PDF, also known as Version of record

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Leonardos, S., Reijnders, D., & Piliouras, G. (2020). PREStO: A Systematic Framework for Blockchain Consensus Protocols. *IEEE Transactions on Engineering Management*, 67(4), 1028-1044.
<https://doi.org/10.1109/TEM.2020.2981286>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

PREStO: A Systematic Framework for Blockchain Consensus Protocols

Stefanos Leonardos , Daniël Reijnsbergen, and Georgios Piliouras

Abstract—The rapid evolution of blockchain technology has brought together stakeholders from fundamentally different backgrounds. The result is a diverse ecosystem, as exemplified by the development of a wide range of different blockchain protocols. This raises questions for decision and policy makers: How do different protocols compare? What are their tradeoffs? Existing efforts to survey the area reveal a fragmented terminology and the lack of a unified framework to reason about the properties of blockchain protocols. In this article, we work toward bridging this gap. We present a five-dimensional design space with a modular structure in which protocols can be compared and understood. Based on these five axes—*optimality, stability, efficiency, robustness, and persistence*—we organize the properties of existing protocols in subcategories of increasing granularity. The result is a dynamic scheme—termed the *PREStO framework*—which aids the interaction between stakeholders of different backgrounds, including managers and investors, and which enables systematic reasoning about blockchain protocols. We illustrate its value by comparing existing protocols and identifying research challenges, hence making a first step toward understanding the blockchain ecosystem through a more comprehensive lens.

Index Terms—Consensus protocols, cryptocurrency, equilibrium, incentives, survey.

I. INTRODUCTION

IN THE seminal Bitcoin paper [129], the pseudonymous S. Nakamoto pioneered the use of *blockchains* as a secure way of maintaining a ledger of currency transfers in a trustless peer-to-peer network. In the ten years since, blockchains have grown [53] to underpin a \$100 billion cryptocurrency market [45]. Meanwhile, their applicability is increasingly understood in a broad range of other contexts [37], e.g., the Internet of Things [69], supply chain management [109], healthcare [126], etc. This rapid growth has induced a considerable number of established market parties to invest in the sector [52], [54], or even develop their own platforms. Noteworthy examples of the

latter include *Quorum* [145], which is developed by JPMorgan Chase, and the *HyperLedger* umbrella project [33], hosted by the Linux Foundation and supported by, *inter alia*, IBM and Intel. Applications of Quorum include JPMorgan's internal digital currency [154] and the Interbank Information Network [39], [134], a platform for cross-border money transfers. Applications of HyperLedger include a project by the U.S. retailer Walmart to track the movement of vegetables [95], [162]. IBM by itself had 1500 employees working on 500 blockchain-related projects in September 2018 [76]. Meanwhile, new multipurpose blockchain platforms developed by startups continue to emerge, e.g., *Ethereum* [30], *Cardano* [35], [104], *Algorand* [26], [82], and *Zilliqa* [182], [183].

This proliferation of blockchain technologies and applications has brought together stakeholders with fundamentally different degrees of technical expertise. So far, the discourse between these groups has been marked by the use of sometimes incongruous terminology, and the lack of a unified communication framework [144]. This hampers the ability of managers and investors to make business decisions, and of newly proposed protocols to be compared and understood. Particularly affected are one of the most fundamental technical aspects of blockchain platforms: the *consensus protocols*.

Consensus protocols fulfill, in a decentralized setting, the role that a single authority has in a centralized database or ledger. It is the mechanism to reach agreement among self-interested peers, and for making consistent decisions out of mutually exclusive alternatives. The choice of consensus protocol has a major impact on a platform's performance, including its security and throughput, and is therefore important for anyone who is involved in blockchain development [121], particularly executives. This can be challenging if the differences between the alternatives are not well-understood.

A. Statement of Contribution

In this article, we address these difficulties by developing an accessible, yet comprehensive framework to improve the communication between the diverse participants of the blockchain ecosystem. We assume only a basic understanding of mathematics and the high-level idea behind blockchains, and introduce technical terms related to blockchains and cryptocurrencies from the bottom up. An extended version of this article, which includes technical definitions of the main concepts, can be found online.¹

Manuscript received June 28, 2019; revised February 14, 2020; accepted March 10, 2020. Date of publication June 28, 2019; date of current version October 9, 2020. This work was supported by the National Research Foundation (NRF), Prime Minister's Office, Singapore, under its National Cybersecurity R&D Program under Award NRF2016NCR-NCR002-028 and administered by the National Cybersecurity R&D Directorate. The work of Georgios Piliouras was supported in part by the SUTD Grant SRG ESD 2015 097, in part MOE AcRF Tier 2 Grant 2016-T2-1-170, and in part by the NRF 2018 Fellowship NRF-NRFF2018-07. Review of this manuscript was arranged by Department Editor K.-K. R. Choo. (Corresponding author: *Stefanos Leonardos*.)

The authors are with the Singapore University of Technology and Design, Singapore 487372, Singapore (e-mail: stefanos_leonardos@sutd.edu.sg; daniel_reijnsbergen@sutd.edu.sg; georgios@sutd.edu.sg).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TEM.2020.2981286

¹[Online]. Available: <https://arxiv.org/abs/1906.06540>




Fig. 1. PREStO framework as a nesting doll of goals.

TABLE I
FIVE AXES OF PREStO AND THEIR MAIN PURPOSE

Dimension	Description	Section
Optimality	Does the protocol maximize the quality of its core outcomes under normal circumstances?	II
Stability	Is the designed protocol an equilibrium?	III
Efficiency	How does the protocol utilize its different resources, e.g., time, space, energy, network bandwidth?	IV
Robustness	Does the protocol's performance withstand perturbations to its parameters?	V
Persistence	If the protocol is forced out of equilibrium, does it recover?	VI

Our main contribution is the *PREStO framework* [42], [178], which is a dynamic tool to identify and classify the properties of blockchain protocols. It is an acronym (in reverse order) of its five main axes: optimality, stability, efficiency, robustness, and persistence, cf. Fig. 1. In Table I, we capture the essence of each category in a single question.

PREStO's modular structure sets it apart from related efforts and enhances its value for managers. Initially, the five categories can be seen as a nesting doll of design goals, where each category considers a wider range of desirable properties than the previous, cf. Fig. 1. We start at the very basic—i.e., optimal performance under ideal conditions—and gradually build up to the more advanced—e.g., recovery mechanisms to survive in the long run. Subsequently, the axes are organized into subcategories of increasing granularity, and PREStO develops into a dynamic tool to identify and group together challenges and research opportunities for the various blockchain protocols, cf. Table IV. We demonstrate its practical use via two running use cases, Bitcoin and Quorum, and conclude with a schematic illustration of the resulting classification in Fig. 5. Furthermore, we extensively draw from the existing literature to motivate our framework.

B. Growing Ecosystem

The consensus protocol introduced by the first blockchain platform—Bitcoin—is commonly called *Nakamoto consensus* [165]. It was designed to work in a *permissionless* setting, i.e., a setting in which any node in the network is allowed to add

data to the blockchain. To prevent network overflow, nodes who seek to extend the blockchain must spend computational effort through a process called *mining*. In the presence of competing chains, honest nodes accept the chain with the most effort spent on creating it. Together, these rules ensure that if more than 50% of the computational power is in the hands of honest parties, then their chain will grow faster than all others. Nodes are compensated for the spent computational power through *rewards* in the form of *tokens* logged on the blockchain. Variations of Nakamoto consensus are currently implemented in over 600 cryptocurrencies [45], [180], including the Ethereum platform and various Bitcoin spin-offs.

In recent years, Nakamoto consensus has increasingly drawn criticism for its *low transaction throughput* and *high energy consumption*. A single Bitcoin transaction costs more energy than 100 000 Visa transactions, and the Bitcoin network as a whole consumes as much energy as a medium-sized country [56]. Furthermore, it is *insecure* in the sense that smaller platforms are vulnerable to attackers who seize a majority of the computational power, as witnessed by the recent 51% attacks on Ethereum Classic [99] and Bitcoin Gold [149]. Finally, research [64], [81], [153], [180] has shown that Nakamoto consensus can be *incentive incompatible*, i.e., participants can increase their rewards by deviating from the protocol. To address these weaknesses, a multitude of new consensus protocols have been proposed that more closely follow traditional theory on *permissioned* (i.e., not open) networks. In particular, many approaches use variations of *Byzantine fault tolerant* (BFT) protocols [112] or other classical consensus protocols such as Paxos [111] and Raft [136]. Such approaches can achieve gains in efficiency and security at the cost of centralization. However, a precise description of this tradeoff is complicated due to the differences between BFT protocol implementations, and the lack of alignment between the terminology used by different parties. This motivates the need for a formal framework to describe and compare different consensus protocols.

C. Related Work

The necessity of developing a unified communication framework for the blockchain ecosystem was already acknowledged in [179]. Accordingly, a brief outline of the PREStO framework was first introduced in [42]. While the five main axes remain the same, their organization into subcategories is first deployed in this article.

The rapid growth of the blockchain-related literature has also stimulated other projects that survey the area from different perspectives. Focusing exclusively on the Bitcoin blockchain, Conti *et al.* [48] provide a systematic review of Bitcoin's underlying features, particularly its security and privacy-related threats and vulnerabilities, and discuss directions for future research. Their analysis extends initial analyses of the backbone protocols of the main cryptocurrencies [72], [73], [75], [138]. In [165], further insight is provided into the development and functionality of the Bitcoin blockchain, in addition to a nonexhaustive, yet interesting timeline of papers related to the analysis of Nakamoto consensus.

In a spirit closer to the present article, Wang *et al.* [173] acknowledge the lack of a comprehensive literature review on the various layers of blockchain technology, and provide a rigorous vision on the organization of blockchain networks. Their work extends to all aspects of the relevant technology and provides a central reference for future work. They define four layers for any blockchain system, from top to bottom: 1) the application layer, 2) the virtual machine layer, 3) the consensus layer, and 4) the network layer. In the present article, we focus on the third (i.e., consensus) layer. That is, application-layer properties, virtual-machine-layer properties (e.g., secure smart contract languages such as Scilla [158]), and network-layer properties (e.g., vulnerability to eclipse [90], BGP hijacking [4], or DoS [101] attacks) are treated only if and when they affect the consensus layer.

The difficulty to conceptualize the dramatically evolving design landscape of blockchains is further supported by [15]. Similar to the present work, they focus on the consensus layer and discuss the various themes and key approaches that are exhibited by current blockchains. They systematize distinctive features and technical properties of existing consensus protocols and provide thorough comparisons, open questions, and directions for future research. Despite the common perspectives, our approach distinguishes itself from [15] due to its mathematical framework that allows for a description of properties from the ground up.

Using a practice-oriented focus, Dinh *et al.* [59] develop BLOCKBENCH, a promising and publicly available software program that is designed to test and compare the performance of blockchain protocols. It applies to private blockchains and its findings are mainly associated with properties in the categories of optimality and efficiency of the PREStO framework, cf. Sections II and IV. The article features use cases of the Ethereum, Parity, and Hyperledger blockchains and concludes that these systems are still far from large-scale adoption. Finally, a nonexhaustive list of related surveys with focal points ranging from smart contract execution to general blockchain applications and research perspectives includes [9], [10], [23], [34], [37], [49], [58], [74], [103], [159], [167], [177], [181].

D. Outline

The rest of this article is organized as follows. In Sections II and VI, we describe the main PREStO axes and define their subcategories. We summarize related issues and open questions related to each category in Section VII. Finally, Section VIII concludes this article.

II. OPTIMALITY

Optimality is the most basic property of a protocol, and generally refers to whether the protocol is optimal within its operational scope. In our setting, it concerns the following question.

Q: *Under normal conditions, does the protocol provide its core functionality in an optimal way?*

By “normal conditions,” we mean that nodes do not act strategically or maliciously, and that there are no capacity constraints. However, we do consider network latency and nodes going

offline. “Core functionality” primarily refers to the functionality of any distributed database, i.e., to correctly read and write to the database. However, some protocols also provide additional functionalities, e.g., a broader notion of transaction types, or a higher level of privacy.

A. Liveness and Safety

Since blockchains are essentially data structures, they must adequately perform the *read* and *write* operations that are required of any database. We focus on the data in the finalized blocks of the chain since nonfinalized blocks can be overturned [160]. The “write” operation then consists of adding a transaction to a finalized block on the chain. The “read” operation consists of observing that a transaction has made it into a finalized block on the chain.

The ability to write and read correctly is formalized through the notions of *liveness* and *safety*. A liveness fault means that a node is unable to write to the blockchain. A safety fault means either that two honest nodes see different results when reading the database, or that a single node sees different results when reading the database at different times.²

In practice, most protocols satisfy these properties only under certain conditions. In particular, most require that the honest nodes control a given fraction of the consensus-critical resources. For example, Bitcoin is safe only if the honest nodes are over 50% strong in terms of processing power. Even then, the property of safety is guaranteed only in a probabilistic sense.³ Quorum is live only if the honest and not permanently offline nodes are at least $\frac{2}{3}$ -strong, and safe only if the adversarial nodes are less than $\frac{2}{3}$ -strong in terms of authority.

During a network partition, protocols can either satisfy liveness or safety, but not both—this is known as the CAP theorem [83]. Different protocols resolve this tradeoff in different ways. Liveness-oriented protocols such as Nakamoto consensus allow the chain to fork by providing an unambiguous rule of how to resolve such forks when the partition ends, e.g., the longest-chain rule. Safety-oriented protocols such as Tendermint [110] and most other BFT protocols [33], [171] require that a (super)majority of participants sign off on each block. This means that during a network partition, at least one of the branches of the chain stops growing. In other settings, different branches of the chain can grow during a fork, but only one of these branches can *finalize* blocks. Examples include a traditional proof-of-work chain with Casper the Friendly Finality Gadget (FFG) as an overlay (“hybrid” Casper) [31], [32]. It is also possible for protocols to guarantee neither liveness nor safety, e.g., Tangaroa [34].

B. Transaction Scope

Some protocols offer fundamentally different types of transactions than others. For example, Bitcoin only supports monetary transactions, which allows for the entire “state” of the

²These two types of safety faults are practically equivalent: if two nodes see different results, then either the network remains permanently forked, or at least one of them will read a different value at some point in the future.

³However, this probability can be made arbitrarily high by increasing the number of confirmations required to make a block final.

system to be described using unspent transaction outputs (UTXOs). However, protocols that support smart contracts (e.g., Ethereum [30]) require that the clients also store the internal variables of the contracts [50]. This may have an impact of efficiency (see also Section IV), both via reduced throughput due to slower transaction processing, and potentially less straightforward scalability (“state sharding” [183]).

C. Privacy

The choice to put data on a blockchain instead of a centralized database has implications for privacy. On one hand, permissionless blockchains such as Bitcoin do not require identity management, thus favoring privacy. On the other hand, the entire history of transactions is publicly accessible, which may allow for deanonymization. In fact, Bitcoin transactions may be better described as pseudonymous than as anonymous [38]. Cryptographic techniques that improve privacy, e.g., zero-knowledge proofs [84] or ring signatures [148], are available, although they may impose additional computational overhead and therefore impact efficiency. Furthermore, usage pattern analysis can lead to user deanonymization even in privacy-minded platforms such as Zcash [102].

III. STABILITY

Since intended behavior cannot be enforced in decentralized settings, one of the core tasks of consensus protocols is to properly incentivize agents to behave appropriately. This will enable the network to reach an outcome that is both stable and desirable. Importantly, stability does not imply optimality; instead, it is concerned with the following question.

Q: Does the protocol incentivize the intended behavior? Is implementing and following-the-protocol the best possible strategy for participating and prospective nodes?

Game theory and traditional economics provide numerous tools to analyze this setting. Yet, as consensus protocols become more elaborate, the incentives and the required stabilizing mechanisms also become more complicated. These issues are discussed separately in the following.

A. Incentive Compatibility

At its core, incentive compatibility entails that it is in the participants’ best interest to follow the protocol, i.e., that the default strategy is a *Nash equilibrium* [130], [170]. In words, a protocol is incentive compatible if *given that* all other nodes follow the protocol, then it is optimal for an entering (or existing) node to also follow the protocol. This definition relies on some assumptions that are not always satisfied in practice. It assumes that first, each node can take as given that all other agents do follow the protocol and second that all agents are *rational*, i.e., utility maximizers. Also, it requires that utility functions are known for each node. Although this set of assumptions may seem restrictive, it is an essential first step in protocol design to establish the stability of a protocol within a vanilla setting. It is within the scope primarily of robustness and to a lesser extent

of persistence to explore what will happen if these assumptions are violated, cf. Sections V and VI.

Based on the above, the task of the blockchain architect is to design the consensus protocol in a way to induce the desired behavior in practice. Differences between intended and observed behavior should be addressed at this point. The theoretical discipline that models and studies such settings is that of *mechanism design* [12], [116], [128]. Applied in the blockchain context, it aims to determine the rules of the protocol in such a way that individual incentives are perfectly aligned with societal goals.

The notion of incentive compatibility can be seen beyond just Nash equilibria. For example, depending on whether the majority is controlled by a single entity or not, one may discern between *strong* and *weak* incentive compatibility [22]. In practice, a consensus protocol of a public, permissionless blockchain needs to properly incentivize rational agents to perform the following actions.

- 1) Participation: acquire protocol resources, e.g., bandwidth.
- 2) Operations: perform core and auxiliary tasks such as proposal and creation of blocks, message propagation, transaction validation and execution, data storage, etc. [38].
- 3) Applications: use the native cryptocurrency or blockchain related applications (“Dapps”).

Although they are integral to viability of existing blockchains, not all of these actions are properly incentivized, and miners’ incentives may be at odds with the underlying protocol [163]. Additional concerns stem from the tension between short-term and long-term incentives [113]. In [146], a consensus protocol is proposed that motivates both ownership and participation, and which aims to develop blockchains for social interaction. In [78], it is shown that the core economic motives for miners—transaction fees and block rewards—are also inherent to the security of Proof of Stake protocols. Apart from the need to incentivize certain operations, like the ones mentioned above, the blockchain protocol also needs to align potentially conflicting incentives of the different entities that are involved in the blockchain ecosystem [68]. Finally, recent works suggest reputation systems as possible solutions to improve the incentive mechanisms of consensus protocols [115], [133].

The theory on social choice and public goods provides insight into misaligned blockchain incentives [155]. A notable instance is captured by the *free-rider* or *pass-the-bucket* problem [17], [164]. In simplified terms, it states that rational agents who benefit from the existence of a public good—in this case, the blockchain—may shift responsibility for its creation to their peers. In the resulting equilibrium, the public good is not created, to everyone’s detriment. In public, permissionless blockchains, this translates to nodes moving costly tasks to other nodes, leading to an improper functionality of the blockchain ecosystem and a deviation from its intended outcome.

A. Impact of protocol resources: Protocol stability is tightly linked to the way that participating nodes acquire and increment their resources, which is starkly different between Proof of Work (PoW) and Proof of Stake (PoS). In PoW protocols such as Bitcoin, computational (CPU) power is the consensus-critical resource. This implies that the costs for participating nodes are mainly electricity and investment in mining equipment [55].

TABLE II
CONCENTRATION OF MINING POWER FOR THE BITCOIN AND ETHEREUM
BLOCKCHAINS (AS OF JUNE 7, 2019) AND CALCULATION OF THE
HERFINDAHL–HIRSCHMAN INDEX (HHI) [147]

BITCOIN			ETHEREUM	
Entity (Pool)	Blocks %		Entity (Pool)	Blocks %
1. BTC.com	20.1%		Ethermine	26.5%
2. AntPool	14.5%		Sparkpool	24.5%
3. F2Pool	13.1%		F2Pool_2	11.8%
4. Slushpool	8.8%		Nanopool	11.2%
5. Poolin	8.8%		MiningPoolHub_1	5.4%
6. ViaBTC	8.3%		Address_1	2.3%
7. BTC.TOP	6.1%		Address_2	1.7%
8. BitFury	4.9%		DwarfPool 1	1.7%
9. BitClub Network	1.7%		zhizhu.top	1.3%
10. Bitcoin.com	1.4%		firepool	1.2%
Total:	87.7%			87.6%
HH Index:	1075.7			1610.5

Sources: blockchain.com and etherscan.io.

PoS protocols generate different dynamics and create different entry barriers. Virtual miners acquire their resources by converting fiat currency to the native cryptocurrency, which they then use as a proof to participate in the consensus mechanism. Mining rewards are again distributed in the native currency, however, in this case, the rewards naturally contribute to the protocol resources. These observations call for a reevaluation of the economics of different protocols through the lens of novel macroeconomic tools [6], [114].

B. Decentralization

Decentralization lies at the core of blockchain design philosophy and is therefore integral for its long-term survival and sustainability [114], [124]. However, existing data show that centralization plagues PoW (and PoS) cryptocurrencies of both high and low market values [25], [71]. Miners join centralized pools that efficiently distribute mining rewards among their participants. This is known to reduce the extreme variance of mining returns that discourages solo miners [150], [156], [166]. Yet, the operation of mining pools introduces unpredictable dynamics in the consensus mechanism and incentivizes miners (or protocol participants) to behave dishonestly, especially under high transaction loads, and destabilize the system [117], [133]. For instance, staking pools—the equivalent of mining pools in PoS protocols—can potentially evolve to become institutions with arbitrary power over their cryptocurrency [25], [65].

Example 1 (Bitcoin and Ethereum): Table II lists the estimated distribution of mining power between the top ten Ethereum mining pools (or accounts) by number of blocks.

The figures indicate a more decentralized market for Bitcoin than for Ethereum. Similar calculations indicate even higher centralization for smaller PoW platforms, for which 51% attacks—distributions in which a single entity owns more than 50% of the resources—are a reality [81], [91], [99]. These figures echo the concerns that the current structure of some of the major blockchain platforms is prone to centralization [4], [5], [80], [114].




Fig. 2. Centralization in the blockchain ecosystem: mining or staking pools, fintech institutions, and intermediaries who offer services—verification, monitoring, data analytics—on public, permissionless blockchains may add value to the ecosystem but also pose a threat to decentralization. The dashed lines indicate possible interconnections between these entities, which may further centralize the system.

From a stability perspective, Nayak *et al.* [131] argue in favor of dispersing mining power since the *follow-the-protocol* strategy ceases to be a Nash equilibrium if a single node becomes too strong. Incentives to derive short-term profits from attacks on mining pools threaten the long-term viability of Bitcoin and negatively impact the Bitcoin ecosystem [113]. Johnson *et al.* [101] show that pool size and computational power are the main criteria when deciding whether to launch a network-level attack against a mining pool. These concerns are not only relevant to Bitcoin, but to other PoW blockchains as well [81].

Ideally, nodes should have no motive to band together at all. However, in Bitcoin, banding together always reduces the reward variance, but when the pools get too strong, trust in the system is undermined and Bitcoins will lose value against other (crypto) currencies. For example, the mining pool GHash.IO was forced to take action to reduce their pool size after they surpassed the 50% mark [66].

Mining pools are not the only threat to decentralization. Other sources involve the underlying network layer, the geographic or economic motives to concentrate mining rigs in countries with low energy cost, and the increasingly sophisticated technology that is required to participate in the block creation process [21], [75], [173]. Cong and He [46] study antitrust policies in Turing-complete blockchains, i.e., blockchains that also support smart contract execution, and argue that although smart contracts mitigate information asymmetries and improve social welfare, they also encourage collusions, and hence generate a threat to decentralization. Various sources of centralization in the blockchain ecosystem are illustrated in Fig. 2.

C. Fairness

An integral element of stability in nonpermissioned protocols is *fairness*, which relies on the premise that participating nodes

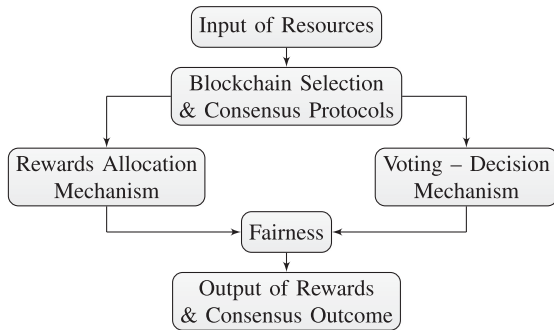


Fig. 3. Fairness in the two core mechanisms of blockchain protocols: allocation of rewards and aggregation of voting data in the block creation process. Rewards are fairly allocated if they are proportional to participants' resources. However, fairness in proportional voting is a premise that has been theoretically challenged [67].

should be rewarded proportionally to their resource contribution [40]. Achieving fairness seems challenging in practice. Message delays and network latency can cause a disproportional distribution of rewards [87]. Focusing on PoS protocols, Fanti *et al.* [65] introduce the notion of equitability to quantify how much a proposer can amplify her stake compared to her initial investment. Consensus in distributed computing with weighted nodes and more general notions of fairness are studied in [7], [77], and [172]. Pass and Shi [139] extend this notion to environments with adaptive corruption by strengthening the definition of “ideal protocol quality” defined in [72] and [138].

Fairness in blockchains can be understood as a two-dimensional notion that entails both the reward allocation and the block creation mechanism, as illustrated in Fig. 3. Current protocols are based on the premise that proportional voting is fair [115]. However, the simple and seemingly appealing axiom “one unit of resource (one computer or one coin), one vote” has been theoretically refuted in traditional voting systems [14], [175].

More importantly, node selection proportionally to their resources – as in current PoW and PoS protocols – does not necessarily imply fairness in the voting process [67]. Interestingly, this also applies to PoW blockchains that do not involve voting. The reasoning is that any coalition that has a combined majority of protocol resources can control the protocol, thus creating a threat in all permissionless blockchains [114].

IV. EFFICIENCY

After establishing that a blockchain protocol is optimal and stable, the next concern is to efficiently meet these goals in practice. The main question in this context relates to a reasonable use of resources and, in particular, to energy waste, scalability of operations, and benchmarking to centralized solutions. Formally, we ask the following question.

Q: *Does the implemented protocol make efficient use of its resources, and how does it compare to a conventional, centralized solution?*

The efficiency of computation, at least from the perspective of time and space, has been thoroughly studied since the 1940s by

Turing and von Neumann. Algorithmic game theory and mechanism design study questions on the intersection of optimality, efficiency, and stability [132]. While, ideally, a protocol implements a (near) optimal equilibrium efficiently, computational complexity theory implies that there are fundamental limitations of efficient computation [137]. These limitations force us to consider tradeoffs, e.g., approximate optimality versus speed or approximation algorithms [168]. The different elements of efficiency are discussed in the following.

A. Energy Consumption

Efficiency exhibits a tradeoff between a modest (excessive) waste of resources and a high (low) risk of attacks. Nodes in PoW protocols provide proofs of the validity of created blocks via energy consumption, which creates a negative environmental externality [108]. A promising alternative is offered by the PoS or *Virtual Mining* protocols, which reduce this huge energy waste [19], [82], [105], [110], [152]. Since PoS protocols delegate decision power via proofs of coin (stake) ownership, one of their main advantages over Nakamoto's PoW is their environmental sustainability [19].

Energy is not the only input that can be inefficiently used by a protocol. Space for data storage, bandwidth, and random access memory is less [41]. Other aspects of efficiency involve the times to process and finalize transactions and the communication complexity that is required for the distributed network to reach consensus [179]. Importantly, different applications introduce various degrees of uncertainty in the use of such resources and increase the challenge of designing efficient solutions. The processing power used for transaction processing and block propagation delay also determines the outcome of the mining competition [119]. Failing to address such issues demotivates agents from participating and leads to centralization. In this sense, efficiency is also related to stability [57].

Better ways to utilize the energy spent in PoW protocols may eliminate—if successful—the advantage of PoS over PoW protocols in terms of energy waste [13]. Vukolić [171] provides a classification of other early proposals and open questions in this direction. Still, all of these alternative proposals need to tackle the problem of *scalability*, described in the following.

B. Scalability

Scalability refers to the property that the consensus protocol—and hence the blockchain—benefits from the addition of nodes or resources [184]. Generally, a blockchain is scalable if it exhibits *positive scale effects*, i.e., if increased participation leads to 1) increased throughput and 2) improved liveness, safety, stability, and efficiency guarantees. Since these indicators may respond differently to variations in the number of nodes (or the amount of resources), it is more convenient to understand scalability as a property of *performance measures* rather than of the blockchain protocol as a whole. Accordingly, a protocol Π is scalable in terms of a performance measure U_{Π} if an increase in the resources of the current state implies an improved performance for U_{Π} . The definition for *negative* performance measures is similar.

Example 2 (Efficiency of Bitcoin): The use of computational resources by the Bitcoin (PoW) protocol to achieve its strong safety guarantees is not efficient [160]: the maximum transaction throughput is the same as five years ago despite a dramatic increase in hashrate and energy consumption [135]. Excessive spending and inefficiencies in the prevailing equilibrium of Bitcoin's *follow-the-protocol* strategy have been identified [20]. It has furthermore been suggested [43] that partial or complete substitution of energy-costly mining activities with PoS mechanisms could benefit Bitcoin and make it more efficient in the long run. Attacks on Bitcoin can inflict a significant energy cost on miners [123]. In general, by partitioning the network or by either censoring or delaying the propagation of blocks, network-layer attacks can cause a significant amount of mining power to be wasted, leading to revenue losses and enabling a wide range of attacks such as double-spending. To deal with these threats, Luu *et al.* [124] propose a mining pool that will run as a smart contract and show that this is a solution with good efficiency and scaling properties.

Currently, a broadly studied solution to scalability is *sharding* (see e.g., Elastico [122], OmniLedger [106], and Ethereum 2.0 [29], [62]). As an alternative approach, Gazi *et al.* [79] model the concept of sidechains as a means to increase scalability and enable the interoperability of blockchains. Their construction features merged-staking, which prevents *Goldfinger attacks*—attacks whose explicit goal is to undermine and destabilize the consensus protocol [21], [108]—and cross-chain certification based on novel cryptographic primitives. Bartoletti *et al.* [16] study a similar combination of consensus protocols with PoS subchains linked to the PoW Bitcoin blockchain.

C. Throughput

Although throughput is closely related to scalability, a protocol can prioritize throughput even without making the protocol fundamentally more scalable. For example, by increasing the maximum number of transactions per block (e.g., Bitcoin Cash vis-à-vis Bitcoin), throughput is increased without essentially affecting scalability. The same is true for protocols such as EOS.IO and TRON, which achieve much higher throughput than, e.g., Bitcoin and Ethereum by curtailing the number of potential block proposers. In fact, a BFT protocol can easily achieve much higher throughput than a Nakamoto protocol if the number of nodes, denoted by N , is low. However, such protocols typically suffer from *negative* rather than positive scale effects when the number of nodes increases due to the $\mathcal{O}(N^2)$ message complexity. So it is possible for a protocol change to have a positive effect on throughput yet a negative effect on scalability.

Fundamentally, *scalability* concerns the effects on the outputs when the resources are changed and the protocol is kept the same, whereas *throughput* (as a PREStO category) concerns the effects on the outputs when the protocol is changed and resources are kept the same.

D. Centralized Systems as Benchmarks

From a managerial perspective, the integral question in launching a blockchain project or application is whether a

blockchain is indeed better than a centralized system for the intended purpose [121], [176]. Since blockchains eliminate trusted authorities to reach consensus via the coordination of distributed and self-interested entities, several questions come into play. How does the distributed system compare to a benchmark solution? Does it provide improved performance in terms of costs, efficiency, and security?

Interestingly, related questions have been thoroughly researched in game theory. In particular, traffic routing, queueing theory and congestion networks explore precisely these tensions between equilibration and efficiency of centrally regulated systems [44]. The tradeoffs are quantified by the *Price of Anarchy* (PoA), which measures the suboptimality caused by self-interested behavior relative to centrally designed and socially optimal outcomes [97], [141], [151]. PoA is defined as the ratio between the performance of the system at the worst case equilibrium and that at a socially optimal state [107].

Studying this question in the current context requires us to quantify different aspects of blockchain performance and compare them to either an existing or a socially optimally (ideal) solution provided by a benevolent social planner or authority. To measure the effects of *decentralizing* a system when implementing it as a blockchain, we evaluate a derivative notion, the *Price of Decentralization* (PoD), which can be defined as

$$\text{PoD}(U_{\Pi}) := U_{\Pi}(D, n)/U_{\Pi}(D, 1). \quad (1)$$

As above, $U_{\Pi} : \Pi \rightarrow \mathbb{R}$ denotes a performance measure of protocol Π . PoD compares the performance of the blockchain at state Π in which the system is operated by n nodes who all follow the protocol, to its performance when it is operated by a single node *and* in an optimal way.

Example 3: To illustrate the above, let $U_{\Pi} : \Pi \rightarrow \mathbb{N}$ be the number of messages that need to be exchanged between N nodes to reach consensus according to protocol Π . In a fully centralized execution of the system, the single entity trivially needs to send one message to each node to inform them of the the decision, leading to $\mathcal{O}(N)$ messages in total. However, in a BFT protocol, in which every node has to send a message to each other node, consensus takes $\mathcal{O}(N^2)$ messages (see e.g., Solidus, Algorand, or Elastico [15]). In this case, $\text{PoD}(U_{\Pi}) = U_{\Pi}(N)/U_{\Pi}(1) = \mathcal{O}(N^2)/\mathcal{O}(N) = \mathcal{O}(N)$. This shows that the PoD of the BFT protocol Π concerning the communication complexity U_{Π} is linear in the number of nodes N and hence, is unbounded as N grows to infinity.

V. ROBUSTNESS

Suppose that a protocol has provable performance guarantees within its scope (optimality) and that the *follow-the-protocol* strategy is an equilibrium (stability), at which the protocol resources are reasonably utilized (efficiency). The next natural step in protocol design is to explore how smoothly and rapidly the protocol's properties degrade when we move away from the vanilla setting. These concerns are expressed by the following question.

Q: *What is the resistance of the protocol to perturbations on its underlying assumptions?*

In the case of a parametrizable protocol, this question may also be phrased in terms of the extent of the parameter variation that the system can tolerate [8]. Essentially, robustness tests the assumptions that were used to equilibrate and stabilize the system. The main challenge is to assess protocol performance under conditions that are not captured by the *ideal* setting of Nash equilibrium, such as parameter fluctuations, collusion between nodes, and malicious or irrational behavior [94].

A. Alternative Equilibrium Concepts

The application of the Nash equilibrium as a stability concept in blockchains is not entirely uncontroversial [11], [88]. In particular, Daskalakis *et al.* [51] and Halpern [88] discuss the following shortcomings of Nash equilibria in distributed computational systems: unexpected behavior (irrational players with out-of-system incentives), coalitional deviations, computational limitations (resource-bounded players), and too much uncertainty or a lack of information (players are unaware of all the aspects of the game). To deal with these issues, Abraham *et al.* [2] and Halpern [88] propose the notion of *robust strategy profiles*, which consist of two defining components: *resilience* and *immunity*.

Despite its theoretical appeal, Halpern [88] observes that the concept of robust equilibria has its own limitations and points to concepts of computational equilibria and particularly to the BAR model—model with Byzantine, Altruistic, and Rational agents—as possible alternatives [3], [11]. Nevertheless, Gradwohl and Reingold [86] provide strong arguments to support the use of Nash equilibria by showing that large games are innately fault tolerant. In fact, *anonymous games* that can be used to model blockchain mining are shown to be resilient against irrational behavior (Byzantine faults), coalitions and asynchronous play.

In an approach that is particularly relevant to the blockchain setting, Liu *et al.* [120] define the robustness of an equilibrium as the maximum proportion of malicious nodes that the desired equilibrium strategy can tolerate, in the sense that this strategy remains the best strategy for rational players. In this definition, robustness is understood as a *local* property, i.e., as a property of a specific strategy profile and against specific adversarial strategies. This definition overcomes the computational difficulties of defining robustness for the blockchain protocol as a whole, and utilizes the fact that in blockchains, the analysis of robustness mainly concerns the *follow-the-protocol* strategy.

B. Out-of-Protocol Incentives

In reality, an adversarial node may try to change the behavior of other nodes by influencing their utility functions through threats or rewards. One of the earliest examples of this is *feather forking* [127] in Bitcoin: in this case, a miner threatens to refuse to extend blocks if they contain a blacklisted transaction. Even if the expected impact of the threat is small, it may be high enough compared to the small cost of enforcing the blacklist to make it rational to comply with the threat. Similarly, *bribery* [21], [22], [125] or *discouragement* [28] attacks can be used to distort the incentives of rational nodes.

In protocols in which it is known how much consensus-critical resources are owned by each of the nodes (i.e., semipermissionless or permissioned blockchains), it may be possible to predict which nodes are scheduled to propose blocks in the near future. Accordingly, Brown-Cohen *et al.* [24] identify two complementary properties—*recency* and *predictability*—of all longest-chain PoS protocols and devise relevant attacks to show that all such protocols are susceptible to certain kinds of malicious behavior. Finally, Seang and Torre [157] explore the tradeoffs between PoW and PoS consensus and find that a combination of both may yield robust results. In particular, for small numbers of participants PoS exhibits better security properties against 51% attacks by mining pools but as the size of the network increases, they recommend reverting to PoW.

C. Resistance to Malicious Behavior

Not all nodes are solely interested in protocol rewards, for example, they may be interested in performing a *Goldfinger attack* [108], in which one cryptocurrency platform is attacked to increase the value of others. One way of modeling this is to assign to such an attacker a utility that is the inverse of the collective utility, and calculate the total losses under the new equilibrium. Another approach is to calculate bounds on the losses that attackers can do relative to their own losses. In [27] and [28], this is made explicit through the *griefing factor*, which is defined as the relative loss that a participating node needs to incur in order to inflict a “unit” of loss on another node.

VI. PERSISTENCE

The four PREStO categories discussed so far consider the protocol when it operates *at or near to* equilibrium conditions. However, what happens if the protocol is forced away from its equilibrium, for instance after a large-scale attack or catastrophic black swan event? Hence, to establish a protocol’s persistence property, we ask the following question.

Q: *Does the protocol have mechanisms to recover from highly nonequilibrium conditions and return to stability in its optimal state? If so, then how fast, and at what cost?*

These questions deal with the long-term sustainability of the blockchain platform. Whereas for robustness, we studied performance under perturbations of the stability assumptions, for persistence, we take this idea to its logical extreme. We assume that the ecosystem is under a large-scale or protracted attack, and study whether it is designed to recover and resume its desirable properties, at least sufficiently often. Hence, we want to assess to what extent a blockchain has the qualities to survive and evolve under extreme crashes, technology shocks or other rare events.

A. Weak and Strong Persistence

To understand protocols from this perspective, we formalize the notions of *weakly* and *strongly* persistent properties in the blockchain context. These ideas have been introduced within the evolutionary game theory and in the study of biological systems, i.e., recovery of an ecosystem after infection from a virus [92], [161]. More relevant to the current context is the




Fig. 4. Dealing with the Blockchain Trilemma: The green dot denotes an ideal protocol that satisfies all three properties in equilibrium. The blue dot denotes a protocol that cycles around the ideal solution and that satisfies the incompatible properties in a weakly persistent (recurrent) manner.

combination of these ideas with tools from optimization theory and algorithm design [18], [142]. Intuitively, a *weakly persistent* property will eventually be satisfied and become satisfied again infinitely often given any initial system condition, whereas a *strongly persistent* property will eventually be satisfied and stay satisfied given any initial system condition [140]. These definitions capture the idea that a desirable property may not be satisfied by a system in equilibrium, but in a dynamic way. This allows for more flexibility between recovery/convergence time, “periodicity,” and the cost of implementation.

Example 4 (The Blockchain Trilemma): As an example, the idea of supporting two or more incompatible properties in a weakly persistent manner, as described above, can be exploited to address the challenging “Blockchain Trilemma” [1], [47], which is illustrated in Fig. 4 and which is discussed in more detail in Section VII-A. The vertices of the triangle correspond to the three seemingly incompatible but desirable properties that blockchain consensus protocols may satisfy: *decentralization*, *scalability*, and *safety*. Protocols can be thought of as points inside the triangle, with coordinates indicating the degree of satisfaction of each of these properties.

Designing an optimal protocol—i.e., a protocol that in equilibrium satisfies simultaneously all three properties—has been a formidable task for blockchain architects [82]. Such a protocol is indicated by the green dot in Fig. 4. However, the idea of weak persistence can be exploited for an alternative design: a protocol could solve the trilemma by constantly alternating between states that satisfy a nonconflicting subset of the otherwise incompatible properties. This is captured by the blue dot protocol and the dashed arrows in Fig. 4, which show its transition between different states.

The idea of studying distributed computation through the lens of dynamical systems has been recently initiated by [98]. Based on their ideas, persistence can be also used to formulate a weaker definition of fairness, cf. Section III-C. Namely, a protocol can be described as *fair* if each node gets to be selected in the block creation process infinitely often.

B. Recovery From Majority Attacks

One of the major challenges in blockchain consensus protocols is the recovery from attacks by malicious nodes who

control the majority of protocol resources [21]. Existing protocols establish their safety and liveness properties under the assumption of either a simple—51%—or an enhanced—usually 67%—honest majority of nodes [181]. Contrary to initial beliefs that these attacks are only of theoretical interests, recent studies have documented the contrary [22]. An important insight from these studies is that it is sufficient to gain control for some short period of time, for instance by temporarily renting protocol resources.

A suggested mechanism to recover from the majority or large-scale attacks on the Ethereum blockchain is the *minority fork*, proposed by [29]. In brief, a minority fork is a mechanism to recover the majority of the consensus-critical resources through a fork initiated by an honest minority. Because the majority cannot create blocks on both branches of the fork, they will be seen as offline on the minority-initiated branch, which may cause their share to shrink on this branch. Such a scheme is fundamentally impossible in permissionless blockchains.

C. Governance and Sustainability

Persistence is closely related to the decision processes that determine the structure and operation of the blockchain. The practical need for an optimal governance structure in the Bitcoin community has already been observed by [108]. In a different approach, Catalini and Gans [38] view the blockchain as a public good and discuss the role of intermediaries that will provide paid services of blockchain verification and monitoring that adds value to the entire blockchain ecosystem. With the exception of some tentative predictions, the formal governance structure of public, permissionless blockchains has yet to be determined [93]. The issues of governance and long-term sustainability in blockchains are integral to their success, and therefore central themes in their evolution.

VII. EVALUATION: USE CASES OF THE PRESTO FRAMEWORK

In this section, we evaluate the PREStO framework’s ability to illustrate the fundamental differences between various protocols. In Section VII-A, we begin by comparing the PREStO framework to the Blockchain Trilemma, another well-known model of protocol properties. Next, we use the PREStO framework to illustrate the properties of a range of nine recently proposed protocols and protocol modifications. As can be seen from our analysis, the PREStO framework reveals more detail than the Blockchain Trilemma. We conclude the section with an overview of research challenges in Section VII-C. Throughout the following section, we refer to the visual summary of the PREStO displayed in Fig. 5.

A. Blockchain Trilemma

In Section VI, we referred to the Blockchain Trilemma as an example of how the theory of dynamical systems can offer a different perspective to one of the most long-standing problems in distributed computing. In particular, largely incompatible properties, like safety, decentralization, and scalability, can be satisfied in a weakly or strongly persistent manner, i.e., in a

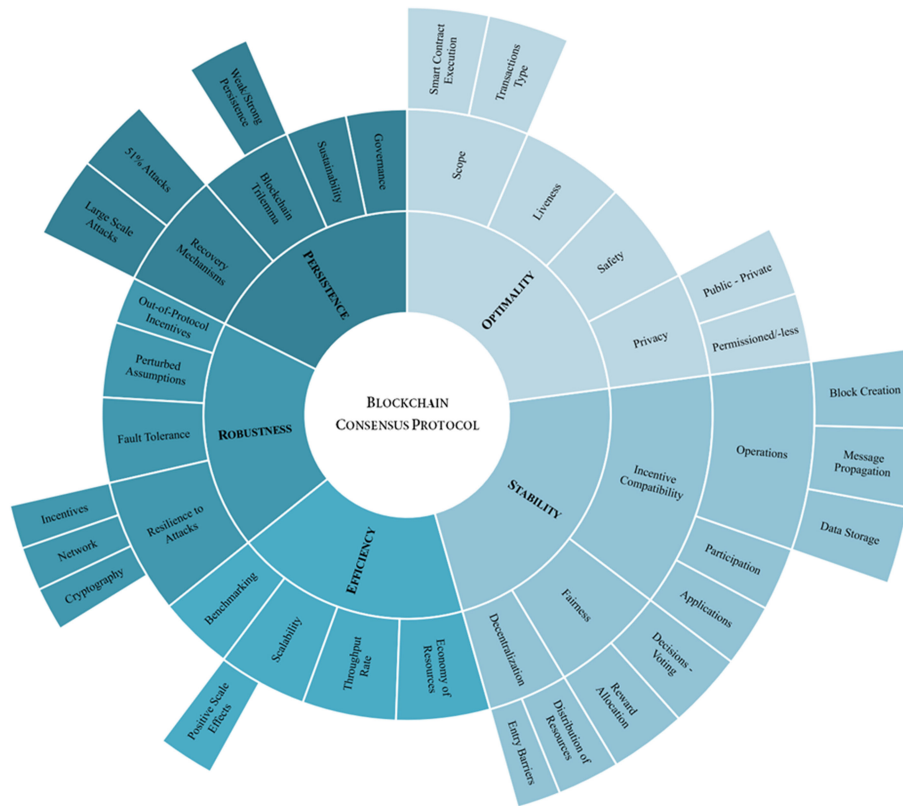


Fig. 5. Visual representation of the PREStO framework. The blockchain consensus protocol lies in the middle of a series of concentric circles. The inner cycle comprises the five major axes of PREStO and the outer cycles correspond to subcategories of increasing granularity. Starting from this setting, the framework can be extended in a dynamic way to integrate features of more elaborate blockchains in the future.

regularly alternating but certainly recurrent manner. This offers a potential way out of the current deadlock in the search for an ideal protocol that will concurrently satisfy all of the three highly desirable properties.

But how do the properties—safety, decentralization, and scalability—of the Blockchain Trilemma relate to the PREStO framework? Or put differently, does PREStO provide the right tool to parse the trilemma in its basic components, to communicate it to experts of different backgrounds and ultimately to reason about and try to resolve it? According to [61], the Blockchain Trilemma describes an asymptotic regime in which certain properties must hold, but leaves these properties open to interpretation. In particular, “security” in the trilemma’s safety property can either refer to liveness or safety of *optimality*, as displayed in Fig. 5. However, a deeper analysis might also consider bribery or discouragement attacks, which are subcategories of *robustness*. The notion of “participating” in the protocol under the trilemma’s decentralization property can refer to the ability to perform the full range of operations, to the existence of entry barriers, or to the ability to be fairly rewarded. All of these are subcategories of the *stability*. Finally, the trilemma’s notion of scalability is clearly related to its namesake in PREStO, which is a subcategory of *efficiency*. However, unlike the trilemma, which requires that throughput is at least linear in terms of the network size, we merely require that it is increasing (one extension of Fig. 5 would be to expand the scalability category to include more subcategories for different rates of growth).

Hence, the PREStO framework can be used to fill the gaps and ambiguities left in the definition of the Blockchain Trilemma. In general, safety relates to optimality (and perhaps robustness), decentralization to stability, and scalability to efficiency. This relationship is illustrated in Fig. 6. Turning to practical examples, Bitcoin is considered to be a safe blockchain because an attacker needs to control at least a given fraction of the total network hashrate for most attack to become possible or profitable (e.g., approximately one third for selfish mining). However, Bitcoin does not scale in its current form, and the emergence of pool-mining has led to an undesirably high centralization [5], [114]—i.e., it is practically impossible for smaller nodes to perform certain protocol operations such as block creation. There is also a tradeoff between decentralization and safety: while solo mining outside of pools leads to a higher degree of decentralization, the large variance of mining rewards may force solo miners to drop out of the mining process due to temporal losses and hence, ultimately decrease the network safety [143].

On the other hand, EOS.IO [60] promises both decentralization and scalability. However, although its minor nodes are able to vote, the ability to create blocks is concentrated in the hands of a small number of nodes. As such, it is unclear if EOS.IO can be described as decentralized. Its relatively small number of active nodes also constitutes a barrier to understanding its exact limits in terms of safety in a large-scale public setting [143]. Similarly, Litecoin and Bitcoin Cash [36], [118] aim to increase scalability—via an increased block frequency

TABLE III
OVERVIEW OF THE IMPACT OF SEVERAL ILLUSTRATIVE PROTOCOL FEATURES ON THE PREStO CATEGORIES

Feature	Example Protocol	Transaction Scope Privacy Decentralization Fairness Scalability Throughput Attack resistance Recoverability						
Partial solutions	<i>FruitChains</i> [139], <i>StrongChain</i> [166]		+	+		-	+	
Smart contracts	<i>Ethereum</i> [30]	+				-	-	
Checkpointing	<i>Casper FFG</i> [31], [32]					-	+	+
Weighted Voting	[115]					+	+	
Zero-knowledge proofs	<i>Zcash</i>	+				-		
Increased block size	<i>Bitcoin Cash</i>		-	-		+		
Increased block frequency	<i>LiteCoin</i>		-	-		+		
Microblocks	<i>Bitcoin-NG</i> [63]		-	+		+		
Sharding	<i>Zilliqa</i> [182], <i>OmniLedger</i> [106]					+	+	-




Fig. 6. Visual representation of the Blockchain Trilemma in relation to the PREStO framework. The tradeoff between Safety, Scalability, and Decentralization is precisely captured by the corresponding subcategories in optimality, efficiency, and stability. Robustness and persistence offer alternative approaches for a long-term resolution of the Trilemma.

or size—but although this increases throughput by a constant factor, the network is still not scalable in the sense that throughput is linear (or even increasing) in terms of the network size. Hyperledger [96] aims to provide both scalability and safety by operating as a private blockchain, i.e., in a less decentralized setting. Algorand [82] and Ethereum [174] are actively researching or developing ideas like proof of stake, sharding, side chains, and more efficient BFT mechanisms that will lead to revolutionary solutions of the Blockchain Trilemma. The list of approaches does not end here, with protocols such as Zilliqa [183], offering yet more ideas to this debate by using sharding to create a scalable but still decentralized protocol. Finally, while many claim that solving the trilemma is essentially not possible, coexistence

of multiple, interoperable blockchains may be another approach to go past this bottleneck in the future [169].

B. Evaluating Features of Blockchain Protocols

In this section, we use PREStO framework as a tool to compare and evaluate a range of recent protocol modifications. A summary of this comparison is presented in Table III, and we present a more detailed discussion given as follows.

1) *Partial Solutions*: Several recent works [139], [166] have proposed to modify Nakamoto consensus by allowing miners to include unsuccessful attempts to solve the PoW puzzle—*partial solutions*—in the blocks. These partial solutions contribute to the block's likelihood to be selected by the fork-choice rule, and award rewards to the finders. Hence, weaker miners are rewarded more often, which reduces their barriers to entry and as such increases *decentralization*. In addition, the inherent advantage that big miners have when confronted with network latency is reduced, which improves *fairness*. Finally, preliminary experiments suggest that it is harder for selfish miners or attackers with a minority of the hash power to overturn blocks, which helps *incentive compatibility* and *attack resistance*. The downside is that new data are added to the blocks, which consumes bandwidth and therefore harms *throughput*.

2) *Smart Contracts*: The main innovation of the Ethereum platform [30] was to extend the functionality of blockchain from token transfers to the creation and execution of Turing-complete programs called *smart contracts*. As such, the *transaction scope* is much wider. This comes at a cost to *throughput* as nodes need to expend considerably more processing power to execute the transactions and update the global state. Finally, smart contract platforms have a more complicated global state than those with just token transfers—this has an impact on the applicability of certain sharding techniques (e.g., transaction versus state sharding) [61], [100] and therefore potentially reduced *scalability*.

3) *Checkpointing*: Casper the FFG [31], [32], a checkpointing protocol for Ethereum, introduces a formal scheme for nodes to create finalized blocks that cannot be overturned without a manual reset. This increases the *attack resistance* of

protocols. Furthermore, Casper FFG enables the “minority fork” mechanism described in Section VI-B, which increase the *recoverability* from majority attacks. However, the voting mechanism that is used for finalization consumes bandwidth, and hence reduces *throughput*.

4) *Weighted Voting*: In [115], a modification to PoS protocols was proposed that weighted the consensus power of nodes not just by their staked deposits, but also by their voting history. For example, if they regularly fail to vote for the blocks on the main chain in a timely manner, then their voting power is reduced. This increases *throughput* because offline nodes are less likely to be elected as block proposers, which means that less time is wasted. By contrast, no new data is added to the blocks and the processing power required for clients to update the voting profiles is negligible. Furthermore, *attack resistance* is improved as attackers who seek to harm liveness by not voting (correctly) rapidly lose consensus power.

5) *Zero-Knowledge Proofs*: In Bitcoin, tokens are protected through a requirement that any token transfer (through the spending of UTXOs) can only be done if correct signatures are produced for the addresses of the sent tokens. These addresses (and the values of individual UTXOs) can be masked using zero-knowledge proofs, which have been implemented in the Zcash platform. This improves the *privacy* of users, but at the cost of additional computing power required to process transactions.

6) *Increased Block Size/Frequency*: By increasing either the size or frequency of blocks (as done in Bitcoin Cash and Litecoin, which are forks/spin-offs of Bitcoin), the *throughput* in terms of the maximum number of transactions per second can be increased. However, this also requires that more resources (especially bandwidth) are consumed per second. Since it is more difficult for more smaller parties to handle this additional load, the barriers to entry are increased, which leads to worse *decentralization*. Furthermore, the impact of network latency is increased, and since larger nodes have an advantage in high-latency situations, *fairness* is reduced.

7) *Microblocks*: The Bitcoin-NG proposal [63] simplifies the leader election process in Bitcoin by dividing time into a sequence of epochs, and keeping the same block proposer for all blocks within the same epoch. This eliminates latency effects within the epoch, and therefore allows for increased *throughput*. The reduction of latency effects also improves *fairness*. However, the smaller number of nodes that participate in node creation harms *decentralization*. Although the ability of slot leaders to perform double-spend attacks within epochs is potentially increased, Bitcoin-NG includes measures to counter this such as “poison transactions” and “proofs-of-fraud,” so the total impact on attack resistance is not entirely clear.

8) *Sharding*: In sharding, the requirement that each node in the network maintains the full transaction ledger is relaxed. Systems that successfully implement sharding have potentially much higher *scalability* and *throughput*. However, in most existing sharding proposals it is much easier to attack a single shard than the entire system, leading to reduced *attack resistance*. Despite some important recent work in this area [106], [182], practical

implementation remains an active research fields with a high potential for future improvements.

In conclusion, the PREStO framework allows for a detailed comparison of different protocols and protocol features that goes far deeper than the safety/decentralization/scalability triad of the Blockchain Trilemma. As we can see from Table III, this allows us to create a “menu” for protocol designers from which they can choose protocol features that complement or offset each other. As the blockchain ecosystem evolves, this table can be both expanded (in terms of protocols) and refined (in terms of subcategories), making the PREStO a truly dynamic tool for the comparison and evaluation of blockchain protocols.

C. Identifying Research Challenges and Opportunities

To provide some insights into possibilities for future work, we summarize the categories and subcategories of PREStO and use them to identify research challenges and opportunities in Table IV. We elaborate on this in the following.

1) *Optimality*: In its current stage, the blockchain ecosystem strives to transition to alternative consensus mechanisms that will retain the success of Nakamoto consensus (which includes PoW) while reducing its energy footprint [89]. Providing formal guarantees of safety and liveness and testing these new consensus mechanisms in large-scale practical settings is an ongoing challenge. The enhanced ability of the next generation of blockchains to enable and secure the widespread execution of smart contracts only adds to the complexity of this already difficult puzzle.

2) *Stability*: A large part of the recent blockchain literature has focused on analyzing the incentives in traditional PoW protocols. However, many questions still remain unanswered. Are (virtual) miners motivated to support the network’s safety during its ups and downs? What are the vulnerabilities—at an incentives level—of the newly proposed PoS protocols and what are the optimal reward schemes that will safeguard their success? Among others, recent advances in theoretical research highlight that centralization and irregular supply of mining power are two threats that are inherent to PoW protocols [5], [70], [85]. Partially, these problems stem from cost asymmetries and economies of scale that manifest in energy-consuming consensus mechanisms. Does PoS remedy these problems and is it indeed the next step in blockchain evolution?

3) *Efficiency*: One pressing—if not the single most important—challenge of the blockchain ecosystem is the issue of scalability. Nakamoto’s PoW proved unexpectedly successful, yet an important hurdle in Bitcoin’s mass adoption is that it does not scale. In particular, even with high energy consumption and increasing environmental externalities, the Bitcoin network cannot process a satisfactory number of transactions per second to compete with the established means of digital payments, like VISA and internet banking. The design of blockchains with scalable properties—in terms of communication overhead, transaction throughput and smart contract execution—remains the cornerstone of research in the Efficiency category.

TABLE IV
CHALLENGES AND CURRENT RESEARCH IN THE DESIGN OF BLOCKCHAIN PROTOCOLS BASED ON THE PREStO FRAMEWORK

PRESTO FRAMEWORK		DESIGN OF BLOCKCHAIN CONSENSUS LAYER Research Challenges – Opportunities
OPTIMALITY	<ul style="list-style-type: none"> • Liveness • Safety • Scope • Privacy features: public/private, permissioned/-less 	<ul style="list-style-type: none"> • Selection of design/architecture & Sybil protection (PoW, PoS etc.). • Exploring the trade-off between safety and liveness. • Secure execution of smart contracts.
STABILITY	<ul style="list-style-type: none"> • Incentive compatibility: Participation, Operations, Applications • Decentralization: Entry barriers, Distribution of resources • Fairness: reward allocation, voting-decision making 	<ul style="list-style-type: none"> • Design of incentive compatible mechanisms. • Protection against adversarial behavior. • Motivate decentralization, fair distribution of resources.
EFFICIENCY	<ul style="list-style-type: none"> • Scalability: positive scale effects • Throughput rate • Economy of resources/ energy consumption • Benchmarking to centralized solutions 	<ul style="list-style-type: none"> • Design of scalable properties. • Reduction of energy footprint. • Compare blockchain to conventional solutions.
ROBUSTNESS	<ul style="list-style-type: none"> • Tolerance of perturbed assumptions/irrational behaviour • Out-of-Protocol Incentives • Resilience to attacks: Incentives – Network – Cryptographic level 	<ul style="list-style-type: none"> • Protection against collusions, Goldfinger attacks. • Equilibration in elaborate adversarial models.
PERSISTENCE	<ul style="list-style-type: none"> • Weak/strong persistent properties • Large scale or majority attacks • Recovery mechanisms: rare events • Governance & long-term sustainability 	<ul style="list-style-type: none"> • Defense against 51% attacks, large network partitions • Blockchain-Trilemma • Design of sustainable blockchains • Decision of governance schemes

4) *Robustness*: To address the problems that arise in terms of robustness, the blockchain ecosystem first needs to establish solutions that satisfy stability and efficiency. The challenges for robustness will push these solutions to their limits and shift the research toward incorporating these systems to every day life. This is where the dynamic nature of PREStO shines most. Given the state of the art when stability and efficiency will have been achieved—if ever—the still open questions and challenges can appear as new categories in the framework's robustness axis. This perspective supports PREStO's current development as a dynamic tool whose purpose is to aid the communication between the increasingly diverse stakeholders of the blockchain community.

5) *Persistence*: The open challenges in this category will also become more relevant at more mature stages in the mass adoption of blockchain technology. With the exact nature of next-generation blockchains and their main applications still unclear, the issues of governance and long-term sustainability will remain at the forefront of blockchain research. Are public blockchains indeed going to support global payment systems and currencies? And if blockchains and cryptocurrencies indeed succeed in this task, what will be the stance of governments, legislators, regulators and policy makers? How will the major technology firms react to the new technology and how will the existing banking system adapt to the changing reality?

The above provides a nonexhaustive list of the open problems that are currently puzzling the blockchain community. Yet, it highlights the comprehensive coverage of the proposed PREStO framework. Owing to its modular structure, PREStO can be expanded or modified to accommodate advances or additional research opportunities in the future of blockchain protocols. Accordingly, it can be used to track the evolution of the blockchain ecosystem and structure the communication between its diverse participants who range from protocol designers, technology experts and end users to academics, corporate managers and strategic investors.

VIII. CONCLUSION

In this article, the PREStO framework saw protocols as multidimensional objects with the following cascade of goals. First, optimality required that the protocol solves the problem that it is defined to address, otherwise there is no good reason to deploy it and the designer should go back to the drawing board. Second, stability aimed to ensure that self-interested agents have an incentive to follow and implement the protocol, i.e., that the protocol itself is an equilibrium. If not, the agents will deviate from it and the deployed protocol will behave unpredictably in practice. Next, efficiency required that resources are used as efficiently as possible (e.g., time, space, network bandwidth, energy, randomness, etc.). Given an optimal, stable, and efficient protocol, the next steps are to consider more elaborate behavioral models from the perspective of the agents. These entail robustness and persistence that measure the resilience of the established equilibria in less idealized settings, and the performance of the blockchain in highly perturbed conditions, respectively.

The exploration of these tradeoffs is an area for multidisciplinary research that relies on the synthesis of ideas from game theory, cryptography, and theoretical computer science. In this direction, PREStO can be used as a dynamic framework to structure the communication between researchers with diverse backgrounds and to accommodate increasingly more elaborate features of future blockchain protocols.

REFERENCES

- [1] J. Abadi and M. Brunnermeier, "Blockchain economics," Nat. Bureau Econ. Res., Cambridge, MA, USA, Working Paper 25407, Dec. 2018.
- [2] I. Abraham, D. Dolev, R. Gonen, and J. Halpern, "Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation," in *Proc. 25th Annu. ACM Symp. Principles Distrib. Comput.*, New York, NY, USA, 2006, pp. 53–62.
- [3] A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J.-P. Martin, and C. Porth, "BAR fault tolerance for cooperative services," in *Proc. 20th ACM Symp. Operating Syst. Principles*, New York, NY, USA, 2005, pp. 45–58.

- [4] M. Apostolaki, A. Zohar, and L. Vanbever, "Hijacking bitcoin: Large-scale network attacks on cryptocurrencies," in *Proc. 38th IEEE Symp. Secur. Privacy*, 2017, pp. 375–392.
- [5] N. Arnosti and S. M. Weinberg, "Bitcoin: A natural oligopoly," in *10th Innovations in Theoretical Computer Science (Leibniz International Proceedings in Informatics)*, vol. 124, A. Blum, Ed. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018, p. 5:1.
- [6] N. Arnosti and S. M. Weinberg, "Bitcoin: A natural oligopoly," in *Proc. 10th Innov. Theor. Comput. Sci. Conf.*, San Diego, CA, USA, Jan. 10–12, 2019, p. 5:1.
- [7] G. Asharov, R. Canetti, and C. Hazay, "Toward a game theoretic view of secure computation," *J. Cryptol.*, vol. 29, no. 4, pp. 879–926, Oct. 2016.
- [8] K. J. Astrom and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton, NJ, USA: Princeton Univ. Press, 2010.
- [9] N. Atzei, M. Bartoletti, and T. Cimoli, "A survey of attacks on Ethereum smart contracts SoK," in *Proc. 6th Int. Conf. Principles Secur. Trust*, New York, NY, USA, 2017, vol. 10204, pp. 164–186.
- [10] S. Azouvi and A. Hicks, "SoK: Tools for game theoretic models of security for cryptocurrencies," 2019, *arXiv:1905.08595*.
- [11] S. Azouvi, A. Hicks, and S. J. Murdoch, "Incentives in security protocols," in *Security Protocols XXVI*, V. Matyáš, P. Švenda, F. Stajano, B. Christianson, and J. Anderson, Eds. Cham, Switzerland: Springer, 2018, pp. 132–141.
- [12] M. Balcan, S. Krehbiel, G. Piliouras, and J. Shin, "Minimally invasive mechanism design: Distributed covering with carefully chosen advice," in *Proc. 51st IEEE Conf. Decis. Control*, 2012, pp. 2690–2695.
- [13] M. Ball, A. Rosen, M. Sabin, and P. N. Vasudevan, "Proofs of useful work," Cryptology ePrint Archive, Rep. 2017/203, 2017. [Online]. Available: <https://eprint.iacr.org/2017/203>
- [14] S. A. Banducci and J. A. Karp, "Perceptions of fairness and support for proportional representation," *Political Behav.*, vol. 21, no. 3, pp. 217–238, 1999.
- [15] S. Bano *et al.*, "SoK: Consensus in the age of blockchains," in *Proc. 1st ACM Conf. Adv. Financial Technol.*, Zurich, Switzerland, 2019, pp. 183–198. <https://doi.org/10.1145/3318041.3355458>
- [16] M. Bartoletti, S. Lande, and A. S. Podda, "A proof-of-stake protocol for consensus on bitcoin subchains," in *Financial Cryptography and Data Security*, M. Brenner *et al.*, Eds. Cham, Switzerland: Springer, 2017, pp. 568–584.
- [17] W. Baumol, *Welfare Economics and the Theory of the State*. Cambridge, MA, USA: Harvard Univ. Press, 1952.
- [18] A. Ben-Tal and A. Nemirovski, *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*. Philadelphia, PA, USA: SIAM, 2001.
- [19] I. Bentov, A. Gabizon, and A. Mizrahi, "Cryptocurrencies without proof of work," in *Financial Cryptography and Data Security*, J. Clark, S. Meiklejohn, P. Y. Ryan, D. Wallach, M. Brenner, and K. Rohloff, Eds. Berlin, Germany: Springer, 2016, pp. 142–157.
- [20] B. Biais, C. Bisière, M. Bouvard, and C. Casamatta, "The blockchain folk theorem," Institut d'Économie Industrielle, Toulouse, France, IDEI Working Papers 873, May 2017.
- [21] J. Bonneau, "Hostile blockchain takeovers (short paper)," in *Proc. 5th IFCA Workshop Bitcoin Blockchain Res.*, 2018, pp. 92–100.
- [22] J. Bonneau, E. Felten, S. Goldfeder, J. Kroll, and A. Narayanan, "Why buy when you can rent? Bribery attacks on bitcoin-style consensus," in *Proc. Financial Cryptography Workshops*, 2016, pp. 19–26.
- [23] J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten, "SoK: Research perspectives and challenges for bitcoin and cryptocurrencies," in *Proc. IEEE Symp. Secur. Privacy*, 2015, pp. 104–121.
- [24] J. Brown-Cohen, A. Narayanan, C.-A. Psomas, and S. M. Weinberg, "Formal barriers to longest-chain proof-of-stake protocols," in *Proc. ACM Conf. Econ. Comput.*, 2019, pp. 459–473.
- [25] L. Brünjes, A. Kiayias, E. Koutsoupias, and A.-P. Stouka, "Reward sharing schemes for stake pools," 2018, *arXiv:1807.11218*.
- [26] Business Wire, "Algorand secures \$62m in funding and announces appointment of executive team," 2018. Accessed on: Feb. 26, 2019. [Online]. Available: <https://www.businesswire.com/news/home/20181024005053/en/Algorand-Secures-62M-Funding-Announces-Appointment-Executive>
- [27] V. Buterin, "The triangle of harm," 2017. Accessed on: Sep. 3, 2018. [Online]. Available: https://vitalik.ca/general/2017/07/16/triangle_of_harm.html
- [28] V. Buterin, "Discouragement attacks," 2018. Accessed on: Jun. 13, 2019. [Online]. Available: <https://github.com/ethereum/research/blob/master/papers/discouragement/discouragement.pdf>
- [29] V. Buterin, "Ethereum 2.0 spec—Casper and sharding," 2018. Accessed on: Oct. 30, 2018. [Online]. Available: <https://github.com/ethereum/eth2.0-specs/blob/master/specs/beacon-chain.md>
- [30] V. Buterin *et al.*, "A next-generation smart contract and decentralized application platform," 2014. Accessed on: Apr. 13, 2018. [Online]. Available: <https://whitepaperdatabase.com/ethereum-eth-whitepaper/>
- [31] V. Buterin and V. Griffith, "Casper the friendly finality gadget," 2017, *arXiv:1710.09437*.
- [32] V. Buterin, D. Reijnders, S. Leonardos, and G. Piliouras, "Incentives in Ethereum's hybrid casper protocol," in *Proc. IEEE Int. Conf. Blockchain Cryptocurrency*, May 2019, pp. 236–244.
- [33] C. Cachin, "Architecture of the hyperledger blockchain fabric," in *Proc. Workshop Distrib. Cryptocurrencies Consensus Ledgers*, vol. 310, pp. 1–4, 2016.
- [34] C. Cachin and M. Vukolić, "Blockchain consensus protocols in the wild," in *31st Int. Symp. Distr. Comput.*, Dagstuhl, Germany, volume 91, 2017, pp. 1:1–1:16, doi:10.4230/LIPIcs.DISC.2017.1.
- [35] C. Cardano. [Online]. Available: <https://www.cardano.org/en/home/>. Accessed on: Feb. 26, 2019.
- [36] B. Cash, "Bitcoin cash: Peer-to-peer electronic cash," 2019. [Online]. Available: <https://www.bitcoincash.org/>. Accessed on: Dec. 15, 2019.
- [37] F. Casino, T. K. Dasaklis, and C. Patsakis, "A systematic literature review of blockchain-based applications: Current status, classification and open issues," *Telematics Inform.*, vol. 36, pp. 55–81, 2019.
- [38] C. Catalini and J. S. Gans, "Some simple economics of the blockchain," Nat. Bureau Econ. Res., Cambridge, MA, USA, Working Paper 22952, Dec. 2016.
- [39] F. Chaparro, "Banks have a big appetite to join JPMorgan's blockchain party," Business Insider, 2018. Accessed on: Feb. 25, 2019. [Online]. Available: <https://www.businessinsider.sg/blockchain-jpmorgan-says-banks-have-big-appetite-to-join-party-2018-2/?r=UK>
- [40] X. Chen, C. Papadimitriou, and T. Roughgarden, "An axiomatic approach to block rewards," in *Proc. 1st ACM Conf. Adv. Financial Technol.*, New York, NY, USA, 2019, pp. 124–131.
- [41] A. Chepurnoy, V. Kharin, and D. Meshkov, "A systematic approach to cryptocurrency fees," Cryptology ePrint Archive, Rep. 2018/078, 2018. [Online]. Available: <https://eprint.iacr.org/2018/078>
- [42] V. Chia *et al.*, "Rethinking blockchain security: Position paper," in *Proc. 1st IEEE Int. Conf. Blockchain*, 2018, pp. 1273–1280.
- [43] J. Chiu and T. Koeppl, "The economics of cryptocurrencies—bitcoin and beyond," School Econ. Finance, Victoria Univ. Wellington, Wellington, New Zealand, Working Paper Series 6688, 2017.
- [44] G. Christodoulou and E. Koutsoupias, "The price of anarchy of finite congestion games," in *Proc. 37th Annu. ACM Symp. Theory Comput.*, 2005, pp. 67–73.
- [45] "Coin MarketCap." [Online]. Available: <https://coinmarketcap.com/>. Accessed on: Feb. 25, 2019.
- [46] L. Cong and Z. He, "Blockchain disruption and smart contracts," Nat. Bureau Econ. Res., Cambridge, MA, USA, Working Paper w24399, Mar. 2018.
- [47] M. Conti, A. Gangwal, and M. Todero, "Blockchain trilemma solver Algorand has dilemma over undecidable messages," in *Proc. 14th Int. Conf. Availability, Reliability and Security*, Canterbury, UK, Jan. 2019, p. 8.
- [48] M. Conti, E. S. Kumar, C. Lal, and S. Ruj, "A survey on security and privacy issues of bitcoin," *IEEE Commun. Surv. Tut.*, vol. 20, no. 4, pp. 3416–3452, Oct.–Dec. 2018.
- [49] K. Croman *et al.*, "On scaling decentralized blockchains (position paper)," in *Financial Cryptography and Data Security*, J. Clark, S. Meiklejohn, P. Y. Ryan, D. Wallach, M. Brenner, and K. Rohloff, Eds., Berlin, Germany: Springer, 2016, pp. 106–125.
- [50] S. Das, A. Kolluri, P. Saxena, and H. Yu, "On the security of blockchain consensus protocols (invited paper)," in *Information Systems Security*, V. Ganapathy, T. Jaeger, and R. Shyamasundar, Eds. Cham, Switzerland: Springer, 2018, pp. 465–480.
- [51] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, "The complexity of computing a Nash equilibrium," in *Proc. 28th Annu. ACM Symp. Theory Comput.*, New York, NY, USA, 2006, pp. 71–78.
- [52] M. del Castillo, "Fidelity launches institutional platform for bitcoin and Ethereum," Forbes, 2018. Accessed on: Feb. 25, 2019. [Online]. Available: <https://www.forbes.com/sites/michaeldelcastillo/>

- 2018/10/15/fidelity-launches-institutional-platform-for-bitcoin-and-ethereum/#4e5a766d93c4
- [53] M. del Castillo, "Despite crypto depression, M&A deals set new record," *Forbes*, 2019. Accessed on: Feb. 25, 2019. [Online]. Available: <https://www.forbes.com/sites/michaeldelcastillo/2019/02/13/despite-crypto-depression-ma-deals-set-new-record/#1d6b65952444>
 - [54] M. del Castillo, "Nasdaq leads \$20 million investment in enterprise blockchain startup Symbiont," *Forbes*, 2019. Accessed on: Feb. 25, 2019. [Online]. Available: <https://www.forbes.com/sites/michaeldelcastillo/2019/01/23/exclusive-nasdaq-leads-20-million-investment-in-enterprise-blockchain-startup-symbiont/#18083f3346d1>
 - [55] S. Dhamal, T. Chahed, W. Ben-Ameur, E. Altman, A. Sunny, and S. Pojary, "A stochastic game framework for analyzing computational investment strategies in distributed computing with application to blockchain mining," 2018, *arXiv:1809.03143*.
 - [56] Digiconomist, "Bitcoin energy consumption index." Accessed on: Sep. 3, 2018. [Online]. Available: <https://digiconomist.net/bitcoin-energy-consumption>
 - [57] N. Dimitri, "Bitcoin mining as a contest," *Ledger*, vol. 2, pp. 31–37, 2017.
 - [58] T. T. A. Dinh, R. Liu, M. Zhang, G. Chen, B. C. Ooi, and J. Wang, "Untangling blockchain: A data processing view of blockchain systems," *IEEE Trans. Knowl. Data Eng.*, vol. 30, no. 7, pp. 1366–1385, Jul. 2018.
 - [59] T. T. A. Dinh, J. Wang, G. Chen, R. Liu, B. C. Ooi, and K.-L. Tan, "BLOCKBENCH: A framework for analyzing private blockchains," in *Proc. ACM Int. Conf. Manage. Data*, New York, NY, USA, 2017, pp. 1085–1100.
 - [60] EOS.IO, "EOSIO strategic vision," 2019. Accessed on: Dec. 15, 2019. [Online]. Available: <https://eos.io/strategic-vision/>
 - [61] Ethereum, "Sharding FAQ," Accessed on: Feb. 13, 2020. [Online]. Available: <https://github.com/ethereum/wiki/wiki/Sharding-FAQ>
 - [62] Ethereum, "Sharding roadmap," Accessed on: Sep. 14, 2018. [Online]. Available: <https://github.com/ethereum/wiki/wiki/Sharding-roadmap>
 - [63] I. Eyal, A. E. Gencer, E. G. Sirer, and R. Van Renesse, "Bitcoin-NG: A scalable blockchain protocol," in *Proc. 13th USENIX Symp. Netw. Syst. Des. Implementation*, 2016, pp. 45–59.
 - [64] I. Eyal and E. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," in *Financial Cryptography and Data Security*, N. Christin and R. Safavi-Naini, Eds. Berlin, Germany: Springer, 2014, pp. 436–454.
 - [65] G. Fanti, L. Kogan, S. Oh, K. Ruan, P. Viswanath, and G. Wang, "Compounding of wealth in proof-of-stake cryptocurrencies," *Cham*, 2018, pp. 42–61.
 - [66] C. Farivar, "Bitcoin pool GHash.io commits to 40% hashrate limit after its 51% breach," 2014. [Online]. Available: <https://arstechnica.com/information-technology/2014/07/bitcoin-pool-ghash-io-commits-to-40-hashrate-limit-after-its-51-breach/>
 - [67] D. Felsenthal and M. Machover, *The Measurement of Voting Power: Theory and Practice, Problems and Paradoxes*. Cheltenham, U.K.: Edward Elgar Publ., 1998.
 - [68] S. Feng, W. Wang, Z. Xiong, D. Niyato, P. Wang, and S. S. Wang, "On cyber risk management of blockchain networks: A game theoretic approach," *Apr.* 2018, p. 1.
 - [69] M. A. Ferrag, M. Derdour, M. Mukherjee, A. Derhab, L. Maglaras, and H. Janicke, "Blockchain technologies for the internet of things: Research issues and challenges," *IEEE Internet Things J.*, vol. 6, no. 2, pp. 2188–2204, Apr. 2019.
 - [70] A. Fiat, A. Karlin, E. Koutsoupias, and C. Papadimitriou, "Energy equilibria in proof-of-work mining," in *Proc. ACM Conf. Econ. Comput.*, New York, NY, USA, 2019, pp. 489–502.
 - [71] B. Fisch, R. Pass, and A. Shelat, "Socially optimal mining pools," in *Web and Internet Economics*, N. R. Devanur and P. Lu, Eds. Cham, Switzerland: Springer, 2017, pp. 205–218.
 - [72] J. Garay, A. Kiayias, and N. Leonardos, "The bitcoin backbone protocol: Analysis and applications," in *Proc. Annu. Int. Conf. Theory Appl. Cryptogr. Techn.*, 2015, pp. 281–310.
 - [73] J. Garay, A. Kiayias, and N. Leonardos, "The bitcoin backbone protocol with chains of variable difficulty," in *Advances in Cryptology*, J. Katz and H. Shacham, Eds. Cham, Switzerland: Springer, 2017, pp. 291–323.
 - [74] J. A. Garay and A. Kiayias, "SoK: A consensus taxonomy in the blockchain era," *IACR Cryptol. ePrint Arch.*, vol. 2018 pp. 284–318, 2018.
 - [75] J. A. Garay, A. Kiayias, N. Leonardos, and G. Panagiotakos, "Bootstrapping the blockchain, with applications to consensus and fast PKI setup," in *Proc. Public Key Cryptography*, 2018, pp. 465–495, doi: 10.1007/978-3-319-76581-5_16.
 - [76] A. Garcia, "IBM is betting big on blockchain technology. Is it worth the risk?" *CNN Business*, 2018. Accessed on: Feb. 25, 2019. [Online]. Available: <https://money.cnn.com/2018/09/06/technology/ibm-blockchain-gamble/index.html>
 - [77] V. K. Garg and J. Bridgman, "The weighted Byzantine Agreement problem," in *Proc. IEEE Int. Parallel Distrib. Process. Symp.*, May 2011, pp. 524–531.
 - [78] P. Gazi, A. Kiayias, and A. Russell, "Stake-bleeding attacks on proof-of-stake blockchains," in *Proc. Crypto Valley Conf. Blockchain Technol.*, 2018, pp. 85–92.
 - [79] P. Gazi, A. Kiayias, and D. Zindros, "Proof-of-stake sidechains," *Cryptology ePrint Archive*, Rep. 2018/1239, 2018. [Online]. Available: <https://eprint.iacr.org/>
 - [80] A. Gervais, G. O. Karame, V. Capkun, and S. Capkun, "Is bitcoin a decentralized currency?" *IEEE Secur. Privacy*, vol. 12, no. 3, pp. 54–60, May/Jun. 2014.
 - [81] 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. ACM SIGSAC Conf. Comput. Commun. Secur.*, New York, NY, USA, 2016, pp. 3–16.
 - [82] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling Byzantine Agreements for cryptocurrencies," in *Proc. 26th Symp. Operating Syst. Principles*, 2017, pp. 51–68.
 - [83] S. Gilbert and N. Lynch, "Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services," *ACM SIGACT News*, vol. 33, no. 2, pp. 51–59, 2002.
 - [84] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM J. Comput.*, vol. 18, no. 1, pp. 186–208, 1989.
 - [85] G. Goren and A. Spiegelman, "Mind the mining," in *Proc. ACM Conf. Econ. Comput.*, New York, NY, USA, 2019, pp. 475–487.
 - [86] R. Gradwohl and O. Reingold, "Fault tolerance in large games," *Games Econ. Behav.*, vol. 86, pp. 438–457, 2014.
 - [87] R. Guerraoui and J. Wang, "On the unfairness of blockchain," *École polytechnique fédérale de Lausanne, Ecublens, Switzerland*, Tech. Rep., 2018.
 - [88] J. Y. Halpern, "Beyond Nash equilibrium: Solution concepts for the 21st century," in *Decision and Game Theory for Security*, J. S. Baras, J. Katz, and E. Altman, Eds. Berlin, Germany: Springer, 2011, pp. 1–3.
 - [89] S. S. Hazari and Q. H. Mahmoud, "Comparative evaluation of consensus mechanisms in cryptocurrencies," *Internet Technol. Lett.*, vol. 2, no. 3, 2019, Art. no. e100.
 - [90] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, "Eclipse attacks on bitcoin's peer-to-peer network," in *Proc. 24th USENIX Secur. Symp.*, Washington, DC, USA, 2015, pp. 129–144.
 - [91] A. Hertig, "Blockchain's once-feared 51% attack is now becoming regular," *Coindesk.com*, 2018. [Online]. Available: <https://www.coindesk.com/blockchains-feared-51-attack-now-becoming-regular>. Accessed on: Feb. 26, 2019.
 - [92] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
 - [93] C. Hoskinson, "Ethereum cofounder says blockchain presents 'Governance Crisis'," 2019. Accessed on: Apr. 11, 2019. [Online]. Available: <http://fortune.com/2019/04/08/ethereum-cofounder-governance-charles-hoskinson/>
 - [94] Z. Hu and J. Zhang, "Toward general robustness evaluation of incentive mechanism against bounded rationality," *IEEE Trans. Comput. Social Syst.*, vol. 5, no. 3, pp. 698–712, Sep. 2018.
 - [95] HyperLedger, "Walmart turns to blockchain (and Hyperledger) to take on food traceability and safety," 2018. Accessed on: Feb. 25, 2019. [Online]. Available: <https://www.hyperledger.org/blog/2019/02/21/walmart-turns-to-blockchain-and-hyperledger-to-take-on-food-traceability-and-safety>
 - [96] HyperLedger, "Hyperledger," 2019. Accessed on: Dec. 15, 2019. [Online]. Available: <https://www.hyperledger.org>
 - [97] N. Immorlica, E. Markakis, and G. Piliouras, "Coalition formation and price of anarchy in cournot oligopolies," in *Workshop on Internet and Network Economics*, A. Saberi, Ed., Berlin, Germany: Springer, 2010, pp. 270–281.
 - [98] A. D. Jaggard, N. Lutz, M. Schapira, and R. N. Wright, "Dynamics at the boundary of game theory and distributed computing," *ACM Trans. Econ. Comput.*, vol. 5, no. 3, pp. 15:1–15:20, 2017.
 - [99] G. Jenkinson, "Ethereum classic 51% attack—The reality of proof-of-work," *Cointelegraph.com*, 2019. Accessed on: Feb. 26, 2019. [Online]. Available: <https://cointelegraph.com/news/ethereum-classic-51-attack-the-reality-of-proof-of-work>

- [100] Y. Jia, "Op Ed: The many faces of sharding for blockchain scalability," Bitcoin Magazine, 2018. [Online]. Available: <https://bitcoinmagazine.com/articles/op-ed-many-faces-sharding-blockchain-scalability>
- [101] B. Johnson, A. Laszka, J. Grossklags, M. Vasek, and T. Moore, "Game-theoretic analysis of DDoS attacks against bitcoin mining pools," in *Financial Cryptography and Data Security*, R. Böhme, M. Brenner, T. Moore, and M. Smith, Eds. Berlin, Germany: Springer, 2014, pp. 72–86.
- [102] G. Kappos, H. Yousaf, M. Maller, and S. Meiklejohn, "An empirical analysis of anonymity in Zcash," in *Proc. 27th USENIX Secur. Symp.*, 2018, pp. 463–477.
- [103] A. Kaushik, A. Choudhary, C. Ektare, D. Thomas, and S. Akram, "Blockchain—Literature survey," in *Proc. 2nd IEEE Int. Conf. Recent Trends Electron., Inf. Commun. Technol.*, 2017, pp. 2145–2148.
- [104] A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: A provably secure proof-of-stake blockchain protocol," in *Proc. Annu. Int. Cryptol. Conf.*, 2017, pp. 357–388.
- [105] S. King and S. Nadal, "PPCoin: Peer-to-peer crypto-currency with proof-of-stake," August 2012.
- [106] E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford, "OmniLedger: A secure, scale-out, decentralized ledger via sharding," in *Proc. IEEE Symp. Secur. Privacy*, 2018, pp. 583–598.
- [107] E. Koutsoupias and C. H. Papadimitriou, "Worst-case equilibria," in *Proc. Annu. Symp. Theor. Aspects Comput. Sci.*, 1999, pp. 404–413.
- [108] J. A. Kroll, I. C. Davey, and E. W. Felten, "The economics of bitcoin mining, or bitcoin in the presence of adversaries," in *Proc. Workshop Econ. Inf. Secur.*, 2013, vol. 2013, p. 11.
- [109] N. Kshetri, "Blockchain's roles in meeting key supply chain management objectives," *Int. J. Inf. Manage.*, vol. 39, pp. 80–89, 2018.
- [110] J. Kwon, "Tendermint: Consensus without mining," 2014.
- [111] L. Lamport *et al.*, "The part-time parliament," *ACM Trans. Comput. Syst.*, vol. 16, pp. 133–169, 1998.
- [112] L. Lamport, R. Shostak, and M. Pease, "The Byzantine General's problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, 1982.
- [113] A. Laszka, B. Johnson, and J. Grossklags, "When bitcoin mining pools run dry," in *Financial Cryptography and Data Security*, M. Brenner, N. Christin, B. Johnson, and K. Rohloff, Eds. Berlin, Germany: Springer, 2015, pp. 63–77.
- [114] N. Leonardos, S. Leonardos, and G. Piliouras, "Oceanic games: Centralization risks and incentives in blockchain mining," in *Proc. 1st Int. Conf. Math. Res. Blockchain Econ.*, Apr. 2019, pp. 183–199, doi: 10.1007/978-3-030-37110-4_13.
- [115] S. Leonardos, D. Reijnders, and G. Piliouras, "Weighted voting on the blockchain: improving consensus in proof of stake protocols," in *Proc. Int. Conf. Blockchain Cryptocurrency*, May 2019, pp. 376–384.
- [116] S. R. Leonard Hurwicz, *Designing Economic Mechanisms*. New York, NY, USA: Cambridge Univ. Press, 2006.
- [117] Y. Lewenberg, Y. Bachrach, Y. Sompolinsky, A. Zohar, and J. S. Rosenschein, "Bitcoin mining pools: A cooperative game theoretic analysis," in *Proc. Int. Conf. Auton. Agents Multiagent Syst.*, Richland, SC, USA, 2015, pp. 919–927.
- [118] Litecoin, "Litecoin," 2019. Accessed on: Dec. 15, 2019. [Online]. Available: <https://litecoin.com/en/>
- [119] X. Liu, W. Wang, D. Niyato, N. Zhao, and P. Wang, "Evolutionary game for mining pool selection in blockchain networks," *IEEE Wireless Commun. Lett.*, vol. 7, no. 5, pp. 760–763, Oct. 2018.
- [120] Y. Liu, J. Zhang, B. An, and S. Sen, "A simulation framework for measuring robustness of incentive mechanisms and its implementation in reputation systems," *Auton. Agents Multi-Agent Syst.*, vol. 30, no. 4, pp. 581–600, Jul. 2016.
- [121] Deloitte LLP, "Blockchain: Enigma, paradox, opportunity," 2019. Accessed on: Apr. 10, 2019. [Online]. Available: <https://www2.deloitte.com/ch/en/pages/innovation/articles/blockchain.html>
- [122] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena, "A secure sharding protocol for open blockchains," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2016, pp. 17–30.
- [123] L. Luu, R. Saha, I. Parameshwaran, P. Saxena, and A. Hobor, "On power splitting games in distributed computation: The case of bitcoin pooled mining," in *Proc. IEEE 28th Comput. Secur. Found. Symp.*, 2015, pp. 397–411.
- [124] L. Luu, Y. Velner, J. Teutsch, and P. Saxena, "SMARTPOOL: Practical decentralized pooled mining," in *Proc. 26th USENIX Conf. Secur. Symp.*, Berkeley, CA, USA, 2017, pp. 1409–1426.
- [125] P. McCorry, A. Hicks, and S. Meiklejohn, "Smart contracts for bribing miners," in *Proc. Int. Conf. Financial Cryptography Data Secur.* 2018, pp. 3–18.
- [126] M. Mettler, "Blockchain technology in healthcare: The revolution starts here," in *Proc. IEEE 18th Int. Conf. e-Health Netw., Appl. Services (Healthcom)*, 2016, pp. 1–3.
- [127] A. Miller, "Feather-forks: Enforcing a blacklist with sub-50% hash power," 2013. [Online]. Available: <https://bitcointalk.org/index.php?topic=312668.0>. Accessed on: Sep. 3, 2018.
- [128] R. B. Myerson, *Game Theory*. Cambridge, MA, USA: Harvard Univ. Press, 2007.
- [129] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008. Accessed on: Nov. 14, 2018. [Online]. Available: <https://bitcoin.org/bitcoin.pdf>
- [130] J. Nash, "Equilibrium points in n-person games," *Proc. Nat. Acad. Sci. USA*, vol. 36, pp. 48–49, 1950.
- [131] K. Nayak, S. Kumar, A. Miller, and E. Shi, "Stubborn mining: Generalizing selfish mining and combining with an eclipse attack," in *Proc. IEEE Eur. Symp. Secur. Privacy*, Mar. 2016, pp. 305–320.
- [132] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. New York, NY, USA: Cambridge Univ. Press, 2007.
- [133] M. Nojoudian, A. Golchubian, and L. Njilla, "Incentivizing blockchain miners to avoid dishonest mining strategies by a reputation-based paradigm," in *Proc. IEEE Sci. Inf. Conf.*, 2017, pp. 1118–1134.
- [134] L. Noonan, "JPMorgan widens blockchain payments to more than 75 banks," Financial Times, 2018. Accessed on: Feb. 25, 2019. [Online]. Available: <https://www.ft.com/content/41bb140e-bc53-11e8-94b2-17176bf93f5>
- [135] K. J. O'Dwyer and D. Malone, "Bitcoin mining and its energy footprint," in *Proc. 25th IET Irish Signals Syst. Conf./China-Ireland Int. Conf. Inf. Commun. Technol.*, Jun. 2014, pp. 280–285.
- [136] D. Ongaro and J. Ousterhout, "In search of an understandable consensus algorithm," in *Proc. USENIX Annu. Techn. Conf.*, 2014, pp. 305–319.
- [137] C. H. Papadimitriou, *Computational Complexity*. Hoboken, NJ, USA: Wiley, 2003.
- [138] R. Pass, L. Seeman, and A. Shelat, "Analysis of the blockchain protocol in asynchronous networks," in *Proc. Annu. Int. Conf. Theory Appl. Cryptograph. Techn.*, 2017, pp. 643–673.
- [139] R. Pass and E. Shi, "FruitChains: A fair blockchain," in *Proc. ACM Symp. Principles Distrib. Comput.*, New York, NY, USA, 2017, pp. 315–324.
- [140] G. Piliouras, C. Nieto-Granda, H. I. Christensen, and J. S. Shamma, "Persistent patterns: Multi-agent learning beyond equilibrium and utility," in *Proc. Int. Conf. Auton. Agents Multi-Agent Syst.*, Richland, SC, USA, 2014, pp. 181–188.
- [141] G. Piliouras, E. Nikolova, and J. S. Shamma, "Risk sensitivity of price of anarchy under uncertainty," *ACM Trans. Econ. Comput.*, vol. 5, no. 1, pp. 5:1–5:27, Oct. 2016.
- [142] G. Piliouras and J. S. Shamma, "Optimization despite chaos: Convex relaxations to complex limit sets via Poincaré recurrence," in *Proc. 25th Annu. ACM-SIAM Symp. Discr. Algorithms*, 2014, pp. 861–873.
- [143] Prasanna, "Blockchain trilemma: Explained," 2019. Accessed on: Dec. 15, 2019. [Online]. Available: <https://cryptoticker.io/en/blockchain-trilemma-explained/>
- [144] "Researchers link realism to blockchain's promise," 2018. Accessed on: Apr. 10, 2019. [Online]. Available: <https://www.princeton.edu/news/2018/12/26/researchers-link-realism-blockchains-promise>
- [145] "Quorum," Accessed on: Feb. 25, 2019. [Online]. Available: <https://www.jpmorgan.com/global/Quorum>
- [146] L. Ren, "Proof of stake velocity: Building the social currency of the digital age," 2014.
- [147] S. Rhoades, "The Herfindahl-Hirschman index," *Federal Reserve Bull.*, vol. 79, pp. 188–189, 1993.
- [148] R. L. Rivest, A. Shamir, and Y. Tauman, "How to leak a secret," in *Proc. Int. Conf. Theory Appl. Cryptol. Inf. Secur.*, 2001, pp. 552–565.
- [149] J. J. Roberts, "Bitcoin spinoff hacked in rare '51% attack'," Fortune, 2018. Accessed on: Mar. 15, 2019. [Online]. Available: <http://fortune.com/2018/05/29/bitcoin-gold-hack/>
- [150] M. Rosenfeld, "Analysis of bitcoin pooled mining reward systems," 2011, *arXiv:1112.4980*.
- [151] T. Roughgarden, "Intrinsic robustness of the price of anarchy," in *Proc. 41st Annu. ACM Symp. Theory Comput.*, Bethesda, MD, USA, May 31–Jun. 2, 2009, pp. 513–522.
- [152] F. A. Saleh, "Blockchain without waste: Proof-of-stake," 2017.
- [153] A. Sapirshstein, Y. Sompolinsky, and A. Zohar, "Optimal selfish mining strategies in bitcoin," in *Financial Cryptography and Data Security*, J. Grossklags and B. Preneel, Eds., Berlin, Germany: Springer, 2017, pp. 515–532.

- [154] A. Saxena, "JPMorgan Chase to create digital coins using blockchain for payments," Reuters, 2019. Accessed on: Feb. 25, 2019. [Online]. Available: <https://www.reuters.com/article/us-jp-morgan-blockchain/jpmorgan-chase-to-create-digital-coins-using-blockchain-for-payments-idUSKCN1Q321P>
- [155] K. M. Schmidt and E. Fehr, "A theory of fairness, competition, and cooperation," *Quart. J. Econ.*, vol. 114, no. 3, pp. 817–868, Aug. 1999.
- [156] O. Schrijvers, J. Bonneau, D. Boneh, and T. Roughgarden, "Incentive compatibility of bitcoin mining pool reward functions," in *Financial Cryptography and Data Security*, J. Grossklags and B. Preneel, Eds. Berlin, Germany: Springer, 2017, pp. 477–498.
- [157] S. Seang and D. Torre, "Proof of work and proof of stake consensus protocols: A blockchain application for local complementary currencies," 2018.
- [158] I. Sergey, A. Kumar, and A. Hobor, "Scilla: A smart contract intermediate-level language," 2018, *arXiv:1801.00687*.
- [159] A. Shahaab, B. Lidgey, C. Hewage, and I. Khan, "Applicability and appropriateness of distributed ledgers consensus protocols in public and private sectors: A systematic review," *IEEE Access*, vol. 7, pp. 43622–43636, 2019.
- [160] E. G. Sirer, "Bitcoin guarantees strong, not eventual, consistency," 2016. Accessed on: Feb. 19, 2019. [Online]. Available: <http://hackingdistributed.com/2016/03/01/bitcoin-guarantees-strong-not-eventual-consistency/>.
- [161] H. L. Smith and H. R. Thieme, *Dynamical Systems and Population Persistence*, vol. 118. Providence, RI, USA: Amer. Math. Soc., 2011.
- [162] M. Smith, "In wake of Romaine E. coli scare, Walmart deploys blockchain to track leafy greens," CNN Business, 2018. Accessed on: Feb. 25, 2019. [Online]. Available: <https://news.walmart.com/2018/09/24/in-wake-of-romaine-e-coli-scare-walmart-deploys-blockchain-to-track-leafy-greens>
- [163] Y. Sompolinsky and A. Zohar, "Bitcoin's underlying incentives," *Commun. ACM*, vol. 61, no. 3, pp. 46–53, 2018.
- [164] M. Steffel, E. F. Williams, and J. Permann-Graham, "Passing the buck: Delegating choices to others to avoid responsibility and blame," *Org. Behav. Human Decis. Processes*, vol. 135, pp. 32–44, 2016.
- [165] N. Stifter, A. Judmayer, P. Schindler, A. Zamyatin, and E. R. Weippl, "Agreement with Satoshi—On the formalization of Nakamoto consensus," *IACR Cryptol. ePrint Arch.*, vol. 2018, pp. 400–416, 2018.
- [166] P. Szalachowski, D. Reijbergen, I. Homoliak, and S. Sun, "StrongChain: Transparent and collaborative proof-of-work consensus," in *Proc. 28th USENIX Secur. Symp.*, 2019, pp. 819–836, doi: 10.5555/3361338.3361395.
- [167] R. Thurimella and Y. Aahlad, "The Hitchhiker's guide to blockchains: A trust based taxonomy," 2018. Accessed on: Nov. 15, 2018. [Online]. Available: <https://wandisco.com/assets/whitepapers/the-hitchhikers-guide-to-blockchains.pdf>
- [168] V. V. Vazirani, *Approximation Algorithms*. New York, NY, USA: Springer, 2013.
- [169] T. Vazz, "The real blockchain trilemma," 2019. Accessed on: Dec. 15, 2019. [Online]. Available: <https://medium.com/coinmonks/the-real-blockchain-trilemma-58824b52fe1d>
- [170] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton, NJ, USA: Princeton Univ. Press, 1944.
- [171] M. Vukolić, "The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication," in *Open Problems in Network Security*, J. Camenisch and D. Kesdoğan, Eds. New York, NY, USA: Springer, 2016, pp. 112–125.
- [172] J. R. Wallrabenstein and C. Clifton, "Equilibrium concepts for rational multiparty computation," in *Decision and Game Theory for Security*, S. K. Das, C. Nita-Rotaru, and M. Kantarcioglu, Eds. Cham, Switzerland: Springer, 2013, pp. 226–245.
- [173] W. Wang *et al.*, "A survey on consensus mechanisms and mining strategy management in blockchain networks," *IEEE Access*, vol. 7, pp. 22328–22370, 2019.
- [174] E. Wiki, "On sharding blockchains," 2019. Accessed on: Dec. 15, 2019. [Online]. Available: <https://github.com/ethereum/wiki/wiki/Sharding-FAQ>
- [175] T. Wong, "An application of game theory to corporate governance," *Omega*, vol. 17, no. 1, pp. 59–67, 1989.
- [176] K. Wüst, and A. Gervais, "Do you need a blockchain?" in *Proc. IEEE Crypto Valley Conf. Blockchain Technol.*, 2018, pp. 45–54.
- [177] J. Yli-Huumo, D. Ko, S. Choi, S. Park, and K. Smolander, "Where is current research on blockchain technology? A systematic review," *PLoS ONE*, vol. 11, no. 10, 2016, Art. no. e0163477.
- [178] V. Zamfir, personal communication, 2017.
- [179] A. Zamyatin, N. Stifter, P. Schindler, E. R. Weippl, and W. J. Knottenbelt, "Flux: Revisiting near blocks for proof-of-work blockchains," *IACR Cryptol. ePrint Arch.*, vol. 2018, pp. 415–428, 2018.
- [180] R. Zhang and B. Preneel, "Lay down the common metrics: Evaluating proof-of-work consensus protocols' security," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2019, pp. 175–192.
- [181] Z. Zheng, S. Xie, X. Chen, and H. Wang, "Blockchain challenges and opportunities: A survey," *Int. J. Web Grid Serv.*, vol. 14, no. 4, pp. 352–375, 2018.
- [182] "Zilliqa." [Online]. Available: <https://zilliqa.com/>. Accessed on: Feb. 26, 2019.
- [183] Zilliqa Team, "The Zilliqa technical whitepaper," 2017. [Online]. Available: <https://docs.zilliqa.com/whitepaper.pdf>. Accessed on: Feb. 26, 2019.
- [184] A. Zohar, "Securing and scaling cryptocurrencies," in *Proc. 26th Int. Joint Conf. Artif. Intell.*, 2017, pp. 5161–5165.