The Distributed Computing Column

by

Seth Gilber

National University of Singapore

seth.gilbert@comp.nus.edu.sg

THE DISTRIBUTED COMPUTING COLUMN

Seth Gilbert
National University of Singapore
seth.gilbert@comp.nus.edu.sg

In this issue of the distributed computing column, Marko Vukolić from Protocol Labs considers the future of decentralized blockchain systems. He argues for the importance of "inclusive" systems, like Bitcoin, where any user can participate. And at the same time, he takes the (perhaps contrarian) position that the issue of energy usage in Bitcoin is likely overstated, with the potential benefits outweighing the costs. Together, these points suggest that Bitcoin is well suited to become the base layer of a large-scale distributed distributed, and he lays out a blueprint for just such a vision. Enjoy this (perhaps provocative) new distributed computing column!

On the Future of Decentralized Computing

Marko Vukolić
Protocol Labs[*]

Abstract

Decentralized systems (e.g., blockchain systems) have the potential to revolutionize financial and payment systems, as well as the internet — for the good of humankind and planet Earth. This position paper aims at justifying this standpoint and at laying out a vision for the future of decentralized computing.

We start by revisiting the definition of decentralized systems, briefly surveying the literature on the taxonomy and different facets of decentralization. We complement existing definitions by proposing *Inclusiveness* as a critical facet. We argue that our notion of Inclusiveness rules out some popular candidate technologies for a "base-level" (or L1) blockchain consensus, namely Proof-of-Stake, from replacing Nakamoto's Proof-of-Work (PoW) as the base consensus technology of decentralized systems.

We further discuss why the high energy consumption of Bitcoin's PoW consensus is not wasteful and why Bitcoin should be embraced as the money of the future. We then argue that future decentralized systems should aim at leveraging the "slow-but-very-secure" PoW consensus of Bitcoin, building systems on top of it rather than trying to replace it. Finally, we propose some open problems for decentralized cloud computing research.

1 Introduction

Decentralized systems are a subset of distributed systems. While a distributed system can loosely be defined as "a collection of independent computers that appears to its users as a single coherent system" [51], a basic definition of a decentralized system requires these independent computers to be controlled by multiple authorities, such that no authority is fully trusted by all [52]. In other words, authorities — or operators controlling computers of a system — are assumed to be potentially malicious, or *Byzantine* [36].

[*] The author is an independent computer scientist and the ConsensusLab Lead at Protocol Labs. Email: *marko@protocol.ai*.

This basic definition is, however, insufficient to capture all nuances of decentralization. For example, this definition is satisfied both by a global-scale open membership (i.e., permissionless) system, such as Bitcoin [43], and by a closed membership (i.e., permissioned) system comprising 4 companies implementing a Byzantine fault-tolerant protocol which tolerates one Byzantine fault. While it is intuitive that the latter system in our example is less decentralized than Bitcoin, it should also be obvious that we need a more fine grained methodology for evaluating the *level of decentralization*. To this end, the first contribution of this paper (Sec. 2) is a brief survey of the literature on taxonomy and different flavors of decentralization. We refine the basic definition of a decentralized system of [52] to identify four main decentralization facets of a distributed system: Resilience, Openness, In-Protocol Incentives and Governance.

This methodology will hopefully give the reader a tool to discern genuinely decentralized projects from others that only give an illusion of decentralization. This seems particularly important given an explosion in the number of "decentralized" cryptocurrency projects. For instance, `coinmarketcap.com` lists more than 12'000 cryptocurrencies, many of which lack a concrete use case and have their token supply and network governance controlled by their respective development teams. These projects piggyback on the rising popularity of tokens that have a genuine use case (e.g., Bitcoin) and can sow confusion and create speculative bubbles.

A minority of these projects are actually motivated by improving the state of the art in decentralized computing. For instance, some projects aim to address performance limitations of Bitcoin (in particular its transaction throughput, currently capped at about 7 transactions per second [53]) or reduce its power consumption. These projects are complicated by fundamental tradeoffs that underlie the design of decentralized systems, in particular the tradeoff between scalability and performance and its very level of decentralization. In other words, it is not easy to scale Bitcoin while retaining its security and decentralization.

In these attempts to improve Bitcoin, there seems to be a confusion about its actual use case. If Bitcoin is seen as a simple payment system, its performance indeed could not answer the demands of the slightly less than 8 billion people on planet Earth. However, Bitcoin network could be used as a final settlement layer while the scaling could happen in a hierarchical approach, in so-called layer 2 [29] and higher-layer protocols.

This leaves a seemingly insurmountable issue related to Bitcoin's power consumption, which today already uses roughly 0.1% of the world's energy production [4]. Proposals aiming at addressing this issue involve changing the base consensus protocol of Bitcoin from Nakamoto's Proof-of-Work [43] to an alternative one. A number of top-ranked cryptocurrencies use or plan to use the approach called *Proof-of-Stake* [2] to be more "green" and energy efficient than Bitcoin.

Unfortunately, Proof-of-Stake is not as open and decentralized as Proof-of-Work and is actually more akin, in its essence, to a closed membership (permissioned) system. While this is intuited in open discussions since the proposal of the Proof-of-Stake idea on the bitcointalk forum in 2011 [2], existing definitions of decentralized systems fail to capture this key difference between Proof-of-Work and Proof-of-Stake. To this end, we propose a new property called *Inclusiveness* (Sec. 3), which refines Openness as previously considered in the literature (and which we cover in Sec. 2). In short, an inclusive system designates a decentralized system which provides *equal opportunities* to its participants. Inclusive systems are a subset of open (permissionless) systems: we show that Nakamoto's Proof-of-Work is inclusive and that systems based on Proof-of-Stake are not. Therefore, Proof-of-Stake cannot be used in Layer 1 in inclusive decentralized systems.

While Bitcoin is inclusive, if it is seen as a simple payment system its energy consumption would indeed be too high a price to pay for this property. However, energy usage or, more generally, the cost of a certain technology, should always be evaluated in the context. We propose to re-evaluate Bitcoin's energy consumption considering a different use case for Bitcoin: that of inclusive decentralized *money* (i.e., "peer-to-peer cash" [43]), of the present and, especially, the future.

If one performs a thought experiment of what would happen if (when) Bitcoin becomes the dominant form of money on Earth, the fears of "excessive", "wasteful" and "useless" power consumption of Bitcoin fade away. We will perform exactly such a thought experiment later in this paper (Sec. 4), thus explaining why Bitcoin's power consumption is not wasteful or excessive, and why it is actually good for planet Earth and mankind.

In short, in Section 4 we make an argument that human behavior in the Bitcoin monetary system is incentivized towards savings and rational spending on things we *need*, with low time preference, encouraging long-term planning, preservation of natural resources and sustainability. We contrast this to the current inflationary fiat monetary system, which inherently promotes spending on things we (think we) *want* and consumerism, high time preference (i.e., focus on short-term profits), and where the entire economy is oriented to "growth", which results in producing and consuming things we often do not need, wasting resources. We will also touch upon the equal opportunities in the Bitcoin monetary system and their impact on human freedom and liberties, contrasting them to inequalities in the current monetary system.

Once we agree that it is good for humankind that the Bitcoin network acts as a backbone for future money and once we embrace this thought, it is interesting to explore how we can use such a very secure decentralized network in use cases beyond money. Some interesting projects pursue use cases different from the use case Bitcoin pursues, e.g., offering decentralized storage of large volumes of data (e.g., Filecoin [6] network which today has more than 10 exabytes (EB), or 10 million TB, of storage capacity). Other projects work towards enabling arbitrary computation in a decentralized network. These *decentralized cloud computing* projects have the potential to, one day, challenge the centralized cloud computing operators that dominate today's internet [27]. Since these systems need to eventually scale to the workloads of today's centralized cloud computing and beyond, they clearly need more efficient consensus hierarchies than the one offered solely by Bitcoin's Proof-of-Work. However, this brings back the decentralization and security challenges of consensus protocols other than Proof-of-Work. The key question is: *can we build more efficient decentralized systems that would benefit from the security of Bitcoin*?

To this end, we propose (in Section 5) to approach the design of the future decentralized internet by leveraging Nakamoto's Proof-of-Work as a secure anchor for critical information needed for secure operation of other, more scalable networks. In a sense, we propose to use Bitcoin network as the backbone of the decentralized internet, helping to secure the rest of it. As an example of such use, we discuss a possible approach in which weights of validators in a Proof-of-Stake network, potentially along with its state checkpoints, are anchored into the Bitcoin blockchain, addressing the critical family of so-called *long-range* attacks on Proof-of-Stake [21]. In this design, to complement Bitcoin as a secure store of state/membership anchors, i.e., hashes of the critical state, we propose optionally using a decentralized content addressable storage system, such as Filecoin/IPFS [15] to resolve those hashes. Finally, we outline some open problems motivating future work.

2 Methodology for Evaluating Decentralization in Distributed Systems

In general, *decentralized systems* can be defined as *a subset of distributed systems where multiple authorities control different components and no authority is fully trusted by all [52]*.

For instance, popular cloud and social networks like Google, Facebook or Twitter, are examples of distributed systems. However, these systems are not decentralized, as each one is controlled by a single authority (company). Note that it is not sufficient for a system to simply have its components controlled by multiple authorities in order to be classified as decentralized — the absence of a single trusted authority is needed, meaning that *any component in a decentralized system can be potentially adversarial [52]*, or *Byzantine* [36].

Beyond the above broad definition of a decentralized system, computer science literature considers multiple *facets* of decentralization in an attempt to characterize its nuances, as well as the differences among decentralized systems (see e.g., [44] for a recent survey). We summarize these into the following *decentralization facets*:

1. **Resilience** of the system to adversarial (Byzantine) behavior of its components, i.e., the authorities that control them, as well as the simple disappearance (also called unavailability) of individual components.

   Resilience itself may apply to different properties of the system, namely *safety* and *liveness* [35, 10]. Informally, a safety property of a system stipulates that "bad things" do not happen and a liveness property stipulates that "good things" do eventually happen (i.e., that the system does not stop making progress).

   For instance, an important liveness property of a blockchain system is *censorship* resistance [26], whereas an important safety property of a blockchain system is *double-spend* resistance [43]. We define these properties later, in the context of Bitcoin, in Section 4.1.

   To quantify Resilience, the scientific literature and engineering practice is typically interested in the minimum number of authorities that the adversary needs to compromise to subvert a key property of the system, such as safety or liveness. In the context of blockchains this number is sometimes referred to as the *Nakamoto coefficient*[1] [48]. Intuitively, the higher the Nakamoto coefficient, the higher the level of decentralization. Per the definition of a decentralized system we adopted [52], the system cannot be deemed decentralized if this number is 1 — i.e., if a single participating authority can compromise a key property of the system. Finally, when evaluating the Nakamoto coefficient, it is important to consider possible business relations or shared control structures among otherwise seemingly independent authorities.

2. **Openness** of the system to new participants. In this sense, a widely-used classification of blockchain systems into *permissioned* and *permissionless* systems (see e.g., [40]) reflects this property. Permissionless systems allow participants to self-elect into the system, whereas permissioned systems rely on an external selection process to be admitted into the system — with the authority to choose [participants] typically residing with an institutional or organizational process [40]. In other words, permissionless systems are *open* to any new participant, whereas permissioned systems are not. Some authors define decentralized systems as only those *in which anyone is able*

[1] Honoring Bitcoin's pseudonymous inventor, Satoshi Nakamoto.

to participate [12], effectively restricting the notion of decentralized systems only to open, permissionless systems. As a general principle, even if we accept permissioned systems as decentralized ones, permissionless systems are to be considered more decentralized than permissioned systems.

Some authors further refine the notion of open, permissionless systems focusing on equality of participants within the system. Karakostas et al. [31] define *egalitarianism* in a rather technically involved way aiming at capturing the proportionality of rewards of participants in blockchains compared to their investment. In a related approach, Fanti et al. [24] define *equitability*, which quantifies how much a participant can amplify her token holdings compared to her initial investment. In the next section (Section 3), we argue that these refinements of Openness are not general enough and define a new refinement of Openness, using the notion of *Inclusiveness*.

Finally, some authors recognize *operational decentralization* as a facet of decentralization related to Openness [44]. Intuitively, operational decentralization aims at capturing hardware requirements for participation in the system — the smaller the hardware requirements, the higher the possibility for anyone to participate in the system and, hence, the higher the level of decentralization. For instance, a system which requires large amounts of storage (e.g., hard disk space) to participate in blockchain A would be deemed more centralized than blockchain B which requires less storage space [44].

3. **In-protocol Incentives** refer to the existence of rewards for protocol participants, paid out to protocol participants in the protocol's *native token*. Incentives are an important facet of decentralized systems [44]: Troncoso et al. [52] argue that the development of adequate incentives is necessary to build a successful decentralized system.

   In general, In-protocol Incentives test the Openness of the system. On the one hand, an open system that provides incentives for participants will attract new participants. On the other hand, a seemingly open system that does not provide In-Protocol Incentives effectively limits its Openness, as new participants have less economic rationale to join the system. Such a system may resort to out-of-protocol incentives, in which case incentives are not governed by system software but by people. Out-of-protocol incentives may involve existing participants establishing business and contractual relations with new participants to motivate them to join the system. This approach resembles and is more common in permissioned networks [11], which, as we discussed, do not satisfy Openness.

   In the context of incentives, wealth distribution across token stakeholders is also considered as an aspect of decentralization [44].

4. **Governance** of the system, focusing on power of human stakeholders to influence and change key rules in the system, e.g., through software updates.

   Several parameters for evaluating the decentralization of governance power have been proposed or discussed in the literature. These include:

   (a) *governance of the infrastructure* [25], or *improvement control* [44], often involving the number of developers contributing to systems' codebase and the number of people contributing to the discussion around the system design [12],

   (b) *existence of a public face* [25], which can be defined as a personality and/or institution that is widely recognized as a spokesperson or a representative of the system.

   (c) *owner control*, measured by examining the total tokens accumulated by the stakeholders in the early adoption period. Depending on the consensus mechanism used, such early tokens may give more power to their owners, causing inequalities and centralization — this is particularly relevant in Proof-of-Stake systems [44].

Finally, some authors [44] consider additional facets of decentralization, including the decentralization at the *network layer*, i.e., pertaining to the decentralization of the network that underlies a distributed system, and the decentralization at the *application layer*, which includes, e.g., the diversity of wallets and exchanges. We acknowledge these decentralization facets that go beyond the core of the system itself, opting to focus on systems proper in this position paper.

3 Inclusiveness in Decentralized Systems

In this section, we define *Inclusiveness*, which refines the notion of Openness defined in the previous section. We argue for Inclusiveness as a key property of decentralized systems and show that Proof-of-Stake systems are not inclusive, in contrast to Proof-of-Work systems.

Inclusiveness is somewhat similar to the notions of *egalitarianism* [31] and *equitability* [24]. Compared to Inclusiveness, these notions are less general as they are defined only for protocols with incentives, practically quantifying the linearity of reward distribution compared to the investment made.

Towards defining *Inclusive* systems, we first define the notion of *Equal Opportunities*.

Definition 1 (Equal Opportunities)**.** *A decentralized system provides* Equal Opportunities *if it satisfies both of the following conditions:*

- *(Resource Symmetry) The system allows any (new or existing) participant Bob to have an equal role in the system as any other existing participant Alice, provided Bob makes the same investment in system resources as Alice. Specifically, this means that if we swap the identities (private/public key pairs) of participants Alice and Bob, the resulting system should be indistinguishable from the original system.*

- *(Genuine Openness) The system cannot reach a state in which it prevents Bob from making such an investment. Specifically, Bob's ability to make this investment must never depend on the permission or actions of either Alice or other participants in the system.*

The first condition of the Equal Opportunities property aims at capturing *resource symmetry*, intuitively capturing equality among new and existing network participants. The motivation behind resource symmetry is to measure if the system gives participants equal power in the system (given that their investment is the same), or if it makes some participants "more equal" than the others, e.g., based on discriminating their identities.

For example, two miners in Bitcoin's Proof-of-Work have an equal role and expected rewards in the system if they contribute the same computing (hashing) power to the system (i.e., if they make the same investment in system resources). On the other hand, swapping identities of a participant Alice, who is part of a permissioned system, and Bob, who is not, yields a system which can be distinguished from the original one. In other words, permissioned systems are not resource symmetric. Moreover, not all permissionless systems are resource symmetric.

The second property of Equal Opportunities aims at refining the Openness property. In principle, an open (permissionless) system could be resource symmetric but prevent, in some state, new participants from making an investment in system resources that would allow them to match the investment of existing participants. Arguably, such a system could not be deemed as *genuinely open*.

Finally, we define *Inclusive* decentralized systems.

Definition 2 (Inclusiveness)**.** *A decentralized system is called* Inclusive *if and only if it satisfies Equal Opportunities.*

It is easy to see that inclusive systems are a subset of open (i.e., permissionless) systems. However, not all open systems are inclusive.

Proof-of-Stake and Proof-of-Work permissionless consensus protocols have fundamentally different implications on the decentralization of the network, which are captured by Genuine Openness. In the following, we show that Proof-of-Stake systems do not satisfy this aspect of Equal Opportunities and, hence, are not Inclusive. We also provide a high-level argument that Proof-of-Work systems satisfy Inclusiveness.

In short, in Proof-of-Stake, “miners” do not expend electrical energy for mining but vote with power proportional to the size of their stake, i.e., holdings in the native token dedicated to voting. This not only implies considerably different economical dynamics compared to Proof-of-Work [24], but may outright lead to violation of Equal Opportunities.

To see this, consider the following simple example of a non-inflationary Proof-of-Stake system, i.e., the one with the non-increasing total supply of a token. If existing miners control more than 50% of the stake in the network and are unwilling to sell their stake to new participants, new participants can never reach the stake of old miners, regardless of the size of their investment. This violates the Genuine Openness aspect of Equal Opportunities and, consequently, Inclusiveness. In this sense, the fact that in Proof-of-Stake existing participants can possibly prevent new participants from meaningfully joining the system evokes similarities between permissioned and Proof-of-Stake permissionless systems.

On the other hand, Bitcoin’s Proof-of-Work satisfies Genuine Openness. Namely, the nature of Proof-of-Work consensus (see Sec. 4.1 for details) does not prevent any participant from making an investment into system resources. In particular, and assuming a free market for computing power, as well as absence of scarcity of computing power and energy, existing participants cannot prevent new participants from entering the system. With innovation in computing (Moore’s law), the computing power of the existing participants actually decays in time compared to the computing power available outside the system, which is free to join the network.

Therefore, we conclude that Bitcoin is an Inclusive system, whereas (non-inflationary) Proof-of-Stake based systems are not. We leave as open the following crypto-economics problem: are there variants of Proof-of-Stake which provably satisfy Equal Opportunities?

4 Why Bitcoin Does Not Waste Energy

In the previous section, we argued that the systems based on Proof-of-Stake are not inclusive, as opposed to those based on Proof-of-Work. Therefore, assuming that we accept Inclusiveness as a necessary aspect of decentralization, it follows that Proof-of-Stake blockchain systems cannot be used as the basis for decentralized systems (i.e., for the so-called layer 1 (L1)).

While Bitcoin and its Proof-of-Work could, technically, be used as the layer 1 of decentralized systems, the seemingly insurmountable issue of its “excessive” energy consumption remains. We propose an argument as to why this is a nonissue and support the claim that the energy consumption of Bitcoin is actually good for humankind and planet Earth, arguing that it is neither wasteful nor excessive. We leave formal modeling and proofs of this argument to future work and the future itself.

Towards making such an argument we need to depart from the narrow domains of computer science and engineering.[2] To help see the “big picture”, we reason about Bitcoin from the angle of other sciences such as sociology, economics and philosophy. We believe that this is, in itself, thought-provoking, as it brings to the spotlight the multi-disciplinary implications of Bitcoin, revealing its intrinsic beauty and ingeniousness. The main impact of the novelty of Bitcoin is to be measured in the spheres of socio-economics and metaphysics, not in computer science.

To this end, later in this section (Sec. 4.3), we will perform a thought experiment in which we will discuss the properties of a prospective world in which Bitcoin becomes the pre-dominant money for mankind, or *unit of account*. For the sake of simplicity, we will perform the thought experiment assuming that Bitcoin becomes the *only* currency in use — we believe that most of our conclusions would hold even if alternative currencies continue to exist, so long as Bitcoin becomes the unit of account. Then, we come back to discussing energy expenditure of Bitcoin (Sec. 4.4), provoking the reader to reconsider if this energy expenditure is a fair price to pay for living in such a world.

In the course of the thought experiment, our economics arguments will mostly be made using simple logical thinking, based on infinite vs. capped money supply, in an attempt to address a large audience. However, for readers who prefer a more structured scholarly approach, where appropriate we will refer to and echo the economics views of the so-called Austrian School of economics, and in particular the thoughts of Friedrich A. Hayek, the 1974 Nobel Memorial Prize laureate.[3]

Before this, in order to make this paper self-contained, we briefly present, in Section 4.1, the background behind Bitcoin and briefly evaluate its decentralization (Sec. 4.2) using the methodology of Sections 2 and 3. A reader familiar with Bitcoin may skip the next section, whereas a reader unfamiliar with Bitcoin is encouraged to also read Nakamoto’s original whitepaper [43].

[2] Provided we do not invoke the Simulation Argument [19] to remain in the realm of computer science.

[3] Specifically, Hayek’s “The Constitution of Liberty: The Definitive Edition” [30], first published in 1960.

4.1 Background on Bitcoin

4.1.1 Bitcoin Basics: Use Case and Monetary Policy

Bitcoin [43] is an open-source peer-to-peer computer network for generating and transferring Bitcoin's native token (bitcoin or "BTC") among the users (peers) of the network. Bitcoin was conceived as an electronic cash network to allow online payments to be sent directly from one party to another without going through a financial institution or any other trusted middleman. This was not possible prior to Bitcoin as all electronic payments required trusted intermediaries, unlike physical, in-person, cash or barter transactions.

On a high-level, in Bitcoin, user Alice wishing to send 1 BTC to another user Bob, digitally signs, using her private cryptographic key, a transaction to transfer 1 BTC from an *address A*, that Alice controls, to *address B* supplied to Alice by user Bob. Alice's private key is cryptographically tied to *address A* (which is basically a cryptographic hash of a corresponding public key). Knowledge of the private key allows Alice to have control over her BTC. As a fundamental principle, whoever controls the private keys corresponding to a given address controls bitcoin pertaining to that address.

The main challenge in such a system is that users do not trust each other. Namely, Alice could attempt to *double-spend* her bitcoin. Consider the following example of a double-spend attempt. Alice signs transaction $tx_{Alice-to-Bob}$ in which she transfers 1 BTC from address A she controls, to Bob's address B. However, she also signs a conflicting transaction $tx_{Alice-to-Alice}$ in which she sends 1 BTC from address A to another address A' that Alice also controls.

Which of these conflicting transactions should be actually taken into account? This is the main technical problem that Bitcoin solves. In a process called *consensus*, peers in the Bitcoin network, without trusting each other, agree on the global, totally ordered *ledger* of all transactions in the system.

In our example, all peers in the Bitcoin network would agree on the relative order between the two conflicting transactions $tx_{Alice-to-Bob}$ and $tx_{Alice-to-Alice}$. The first transaction in that order would be considered valid, whereas the other would be discarded.

Besides preventing double-spends (as a safety property), another important property Bitcoin provides is censorship resistance (as a liveness property). In short, censorship-resistance guarantees that a correctly-behaving user Alice will have her transactions eventually included in the ledger (while possibly having Alice pay a *transaction fee* for this service). In other words, censorship-resilience guarantees that transactions will not be excluded from the Bitcoin ledger due to actions of the Byzantine adversary or peers disappearing from the system.

For efficiency reasons, Bitcoin processes transactions in blocks, grouping a number of transactions together, up to a certain maximum block size. Effectively, the Bitcoin consensus mechanism establishes a global order on those blocks forming a *chain* of blocks (i.e., a "blockchain"). Consequently, Bitcoin establishes global order on the transactions contained in those blocks.

Bitcoin software defines a so-called *genesis* block, the first block in the chain, to which the latter blocks are appended. The Bitcoin genesis block contains a link to the "real" (physical) world, by embedding the headline of the cover page of *The Times* (British daily national newspaper) from January 3rd, 2009 reading *"Chancellor on Brink of Second Bailout for Banks"*. This link to the real world, beyond possibly conveying a motivation for the existence of Bitcoin, is important for proving that the creator of the Bitcoin network could not have ran the code prior to January 3, 2009.

At the beginning of the Bitcoin blockchain history, there weren't any bitcoin to transact, as none had been brought into existence (i.e, *minted* or *mined*) yet. To bring bitcoin into existence, Bitcoin software defines a *block reward*, which is an incentive for network participants to take part in Bitcoin consensus. Bitcoin rewards every participant who successfully adds a block to the blockchain with a fixed reward, which halves every 210,000 blocks. The period of 210,000 blocks corresponds roughly to 4 years, as Bitcoin block production time is set to self-adjust to 10 minutes per block on expectation. For the first 210,000 blocks, the block reward was 50 BTC per block. With maximum supply, as stipulated by Bitcoin code, being 21 million BTC, 50% of all bitcoin have been mined in the first 210,000 blocks. With block reward halving to 25 BTC, from block 210,001 to block 420,000, an additional 25% of the total supply have been minted in that period, and so on, with the current Bitcoin block reward conveniently conveying which percentage of the total supply has been minted within the current 4-year window. Currently, more than 12 years after the genesis block, the Bitcoin network has produced over 700,000 blocks with the current block reward being 6.25 BTC.[4] The last 10% of Bitcoin's total supply is to be mined between today and the year 2140.

A participant in the Bitcoin network is an entity that runs a *full node*. Each Bitcoin full node keeps the entire history of the blockchain, validates new blocks and (optionally) participates in creating new blocks. Bitcoin's maximum block size and a relatively conservative time period interval of 10 minutes between blocks imply that the blockchain does not grow too fast compared to advances in computer hardware. Users can opt-out from running full nodes, by maintaining only *client* wallets, which protect their private keys and send Bitcoin transactions to others' (full) nodes.

In the following, we explain how blocks are generated and validated in the

[4] Each bitcoin is divisible into 100 million smaller units, usually called satoshis.

Bitcoin consensus.

4.1.2 Bitcoin Consensus

Bitcoin consensus proceeds as follows [43]:

1. New transactions are broadcast to all nodes.

2. Each node collects new transactions into a block. A node cryptographically links the new block to its predecessor (parent) block. These parent links define the position of the new block in the blockchain, all the way to the genesis block. In short, a node chooses the predecessor block for the new block to be the one which has the longest chain[5] to the genesis block, out of all blocks known to a node. In principle, nodes consider the transaction history defined by the longest chain as the only valid one.

3. In the process often called *mining*, or *Proof-of-Work* [43], each node works on finding a final piece of information, called a *nonce* that, when embedded into the new block, will make other nodes accept and declare the new block as *valid*.

   This is the key point in the otherwise relatively straightforward Bitcoin consensus. Namely, Bitcoin requires a cryptographic hash of a valid block to start with a number of zeros (0s) when represented as a bit string. Since the output of a cryptographic hash function cannot effectively be predicted, a hash of a block with one specific nonce appears basically as a random string of 0s and 1s. Therefore, nodes need to try many nonces in order to be lucky and construct the required final data for the block such that the hash of the block will start with many 0s, as required by the validation code. The number of zeros required is self-adjusted by the network during its lifetime, based on code, to maintain an expected block time of 10 minutes between the blocks.

   Finding a nonce which makes the block valid is effectively a very simple but computationally intensive guessing game in which a node repeatedly tries different nonces, applies them to the rest of the block, applies the hash function and sees if the output hash has the required number of leading zeros.

4. When a node finds a nonce and completes the Proof-of-Work, it broadcasts the block to all nodes.

[5] In fact, it is the chain which requires most work, which is most often the longest chain. For simplicity of narrative, we talk about “longest chain”.

5. Other nodes accept the block only if: *(i)* all transactions in it are valid and do not contain already spent bitcoin, and *(ii)* the hash of the block starts with the required number of 0s. Unlike the mining step (step 3), this validation step (step 5) is very simple and cheap to compute.

In the recent months, the Bitcoin network as a whole is estimated to have performed anywhere between 68 EH/s (exahashes per second) on June 28, 2021 and 190 EH/s (on May 9, 2021). An exahash per second is one quintillion (a billion billion) hashes per second, a very large number of operations. With each hash operation costing actual physical-world energy to compute, this results in large power consumption of the Bitcoin network, which is sometimes frowned upon and considered as wasted.

4.2 Bitcoin's Decentralization

In evaluating Bitcoin's Resilience, we look at two major possible issues: the double-spending (safety) issue and the censorship of transactions (liveness) issue. To mount these attacks effectively on the Bitcoin network, the attacker needs to control more than 50% of the network computing power. This would allow the attacker to simply ignore blocks produced by the rest of the network and produce the dominant longest chain, which would then, by Step 2 of the Bitcoin consensus protocol (Sec. 4.1.2) be the effective history of transactions. In the case of censorship attacks - this new history could simply be empty of transactions. This is known as a 51% attack for Bitcoin and requires a majority of the hash power of the network.

Whereas it is difficult to precisely calculate the Nakomoto coefficient (number of different authorities required to mount the attack) for Bitcoin, this resilience can be (very) conservatively estimated. Namely, to spread out more evenly, in time, their earnings from block rewards, Bitcoin nodes often group into so-called *mining pools*. While individual nodes are often not directly under the control of a mining pool's operator authority and could leave the mining pool if they detected that they were participating in an attack, for a *very conservative* estimate of Resilience one could theoretically assume that a mining pool fully controls all the nodes within. With this in mind, at the time of writing, more than 50% of Bitcoin mining power is controlled by 4 mining pools.

However, in practice, and as indicated in the Bitcoin whitepaper [43], the economic incentives of Bitcoin make safety attacks towards compromising Resilience less likely than if the In-Protocol Incentives did not exist. If certain nodes control a large amount of computing power, they have an economic dilemma between using that power to attack the system or using that power to behave correctly and earn block rewards and transaction fees. This intuitively contributes to effectively increasing the Nakamoto coefficient (Resilience measure) and consequently increasing the decentralization level of the network, in presence of economically rational participants.

As for other facets of decentralization we defined earlier in the paper, Bitcoin is Inclusive (Sec. 3) and has in-Protocol Incentives (Sec. 4.1).

It also has excellent operational decentralization, i.e., has low hardware requirements for running a full node. Today, the size of the Bitcoin blockchain is about 400GB of data, which means that a full node can be easily run on low-cost hardware, with a mid-sized hard-disk (or a larger microSD card) and internet connection, basically by anyone.[6] This last feature is an important decentralization aspect, and has many benefits. We will illustrate one such benefit later, in Section 5.2, outlining a design that requires a user of a Proof-of-Stake blockchain to run a Bitcoin node to prevent a number of critical attacks on Proof-of-Stake. If running a Bitcoin full node would have had high hardware requirements, this approach would be unrealistic.

Low hardware requirements for running a full node do not imply that everyone can be successful in mining. Full nodes are incentivized to invest more into hardware and computing equipment if they wish to have a higher probability of obtaining block rewards in the Bitcoin consensus. It is well known that economically viable Bitcoin mining requires larger investments, with large miners even running datacenter-size operations. Here, we should not confuse equal opportunities and inclusiveness with linear payouts, i.e., rewards proportional to an investment. Like practically any other economic undertaking, Bitcoin mining benefits from economies of scale. This does not undermine its inclusiveness.

Finally, Bitcoin's Governance benefits from the absence of any single individual or company acting as its public face (as Satoshi Nakamoto disappeared from the public discourse more than ten years ago). Regarding owner control, Bitcoin did not have a hidden owner accumulation phase. The first transaction in the Bitcoin network happened in block #170, seemingly between Satoshi Nakamoto and a cryptographer Hal Finney, on January 12, 2009, nine days after The Times newspaper timestamp contained in the genesis block. The first block following the genesis block was mined, probably by Satoshi Nakamoto, six days after the genesis block, on January 9, 2009.[7]

Concerning code improvement proposals, anyone can propose a change to

[6] Bitcoin full node can be run on hardware which today costs about $200 USD, see https://getumbrel.com.

[7] As it is widely believed, Satoshi Nakamoto may have mined a sizeable number of bitcoin in the early days of the network following the genesis, as an early participant. The exact number is practically impossible to support with hard evidence. However, we do have hard evidence, in the very Bitcoin transaction history, that an overwhelming majority of those early bitcoin that could be attributed to Satoshi Nakamoto were never transacted on the network.

the Bitcoin open-source software via Bitcoin Improvement Proposals (BIPs).[8] In practice, relatively few "core" developers (developers of the Bitcoin Core reference node software) propose and implement changes [12]. Major changes to software are relatively rare, with no BIP containing a backwards incompatible change to Bitcoin consensus (also known as a hard fork) ever having been deployed in the software. For changes that restrict consensus validation rules even further and are backwards compatible (soft forks), consensus among core developers is required, together with approval of miners through on-chain voting. That said, as Bitcoin is open-source software, anyone can make any change to the software. Whoever makes such a change, changing e.g., Bitcoin parameters (block size, or frequency, or token supply), has "only" to convince other users to migrate to using such a network. A number of such backwards incompatible changes to Bitcoin code have resulted in Bitcoin network forks (examples include Bitcoin Cash and Bitcoin Gold), all considerably less popular than Bitcoin.

4.3 Life on Earth with Bitcoin as Money and Unit-of-Account: a Thought Experiment

Towards our thought experiment, we first review the incentives for participants (people and businesses) in the current monetary system (Sec. 4.3.1) and then move to incentives for people in the Bitcoin system (Sec. 4.3.2). We then consider the impact of Bitcoin on economic inequalities in Section 4.3.3.

4.3.1 Human Incentives under the Current Fiat Monetary System

As Hayek put it in 1960, *"With government in control of monetary policy, the chief threat in this field has become inflation."* [30]. Hayek made this statement even before we completely abandoned the gold standard in 1971 and transitioned to so-called *fiat* money (not backed by anything), and well before progressive reduction and subsequent elimination (in some jurisdictions) of required reserves for commercial banks when making loans in fiat money.

In such an inflationary economic environment, with continuous increase in monetary supply there are several issues, out of which we outline just a few:

- Savings are silently taxed by inflation, therefore incentivizing people *to spend* fiat money quickly, or to invest it.

- Investments are done through e.g., stocks, by entrusting funds to other economical players, i.e., businesses. The success of these businesses is measured by their economic "growth", where this "growth" is measured again

[8] https://github.com/bitcoin/bips

in the same inflationary fiat currency. This incentivizes businesses to promote spending, which business predominantly do by relying on marketing and ads to push products to people in an attempt to make them spend more.

- Consequently, the *time preference* of both individuals and businesses is high (i.e., they are *short-term* oriented). Individuals are incentivized to spend money quickly, whereas businesses are incentivized to "grow" by pushing products to people, even if these are not *needed* by people. Marketing divisions exist to *create the demand* for products.

Therefore, current inflationary fiat economy encourages and incentivizes spending of resources on otherwise unneeded products, for which demand is simply created to fulfill the goal of selling more goods. It is not difficult to see that, this being the predominant economic model for 8 billion people, is unsustainable and will lead, in the long term, to overexploitation and pollution of the planet, without necessarily improving quality of life. The current economic model lacks incentives for long-term considerations.

This system has other profound issues related to the structure of the monetary system and positioning of the preferred players (e.g., banks) in such a system.[9] As inflation creates more money supply, these preferred players are close to the source of money (i.e., to central banks), creating the so-called Cantillon effect[10] and inequalities in money distribution. These inequalities have profound implications on the very freedom and liberty of people who are at the opposite side of this inequality.

4.3.2 Human Incentives under the Bitcoin Monetary System

The idea of using *"rules versus authorities in monetary policy"* has been argued since, at least, Henry Simons [47] in 1936. As Hayek writes, *"arguments advanced [by Simons] in favor of strict rules are so strong that the issue is now largely on of how far it is practically possible to tie down monetary authority by appropriate rules"* [30].

Bitcoin, for the first time, offers a monetary network with such "strict rules", which cannot be changed or manipulated even by the strongest adversaries. Prior to Bitcoin we, as a humanity, did not have such a technology. Therefore, it is pardonable that we resorted to inferior solutions, including the current fiat monetary

[9] Stretching our decentralization terminology of Sections 2 and 3, the fiat monetary system does not provide Openness and, consequently, is not Inclusive.

[10] Richard Cantillon, an 18th century economist, suggested that inflation occurs gradually, originating the concept of non-neutral money, positing that the original recipients of new money enjoy higher standards of living at the expense of later recipients. [8]

system. However, since we now have this technology, let's explore what kind of the world it promises.

With the total money supply of 2.1 quadrillion monetary units (21M BTC × 100M satoshis/BTC), Bitcoin is a monetary system with hard-capped money supply. This cap cannot be changed in an undetectable way: a single independent full node running Bitcoin software needs not to trust anyone to be able to verify that the supply did not change. The same goes for the Bitcoin's rate of increase in money supply.

In the current adoption period, Bitcoin encourages *long-term savings*. This argument can easily be made without resorting to the historical exchange rate with respect to legacy fiat currencies (although this historical exchange rate can be used to verify the argument). Namely, as adoption grows and more people accept Bitcoin, its value grows, along with its network effect.

While people save and hold (colloquially, *hodl* [3]) Bitcoin, especially for the long term, they change their time-preference and move away from the spending economy and its incentives (Sec. 4.3.1) towards the mindset and an economy which encourages savings and, consequently, preservation of natural resources. While doing this, people have an option to retain financial sovereignty (which can be directly related to personal freedom and liberties [30]) by securing their private keys themselves. It is worth noting that this aspect of Bitcoin still requires a considerable level of technological literacy.

Fast forward to the future in which Bitcoin is the only money for humankind. People do spend their bitcoin,[11] but primarily on things they *need*, as they know the Bitcoin are scarce. It is certainly possible to earn more Bitcoin by working, yet this will entail providing genuine value to other people, in order for the latter to be ready to spend their scarce monetary asset.

This marks a shift from a spending economy in which depreciating money is spent on things we (think we) want but often do not need, to a saving economy of appreciating money which is spent on goods we need (regardless of an individual's definition of a need — this is purely a *local* definition). Business models dramatically change. Classical economic "growth", along with short-term orientation to revenues and profits, becomes largely meaningless under the Bitcoin monetary standard, as one cannot infinitely grow their business when measured in Bitcoin, as it is hard-capped. Instead, businesses need to focus on providing only true value to people and can plan in long-term, since their monetary power does not depreciate in time.

Under the Bitcoin monetary standard, humans as a species are incentivized towards a savings economy and mindset, putting focus on their materialistic needs

[11] This of course occurs already today, yet people who can support themselves on fiat income do not yet need to do this.

instead on their materialistic wants, which were often not even their genuine wants, but were pushed to them in the fiat system. As roughly 8-10 billion people (of the present and future) transition from the fiat monetary system to Bitcoin, the effects on the environment become profound, with resources saved, helping sustainability, *even if this was not necessarily the original goal of every single individual using Bitcoin.*

As their Bitcoin savings allows people to accumulate and preserve the monetary power that they gained by their past work, without fear of an external authority who could depreciate its worth by external decisions, people are now free to decide: do they *want to work* while providing the actual value to others, or do they want to dedicate their time to art, poetry, science, new inventions, charity, or spiritual development, etc? Here, we assume that the reader accepts that we live in an era in which most jobs needed for human basic needs, can be done by the machines and that we are steadily approaching a post-scarcity society, in which all people can be ubiquitously provided basic needs (per Maslow's hierarchy of human needs). Technology already started to free people from *the need to work* and will continue to do so, but only if the distribution of the benefits of such a technology are shared among people.

The success of organizations will primarily be measured not by their revenue and profits, but by their network effect and the number of human lives that they qualitatively improve. An example is the Bitcoin network itself. It does not have a classical business model in the context of the existing monetary system, nor does it need one. It changes lives by pure adoption, giving people back their *freedom* and their *time*. We can only imagine the potential of a society of free people, who are not coerced to trade their time, the only scarce resource we humans actually have, for ever-inflating money, and who are incentivized by the economic system to save and mind their spending. In this sense, it is well probable that Bitcoin will help us to evolve as a species.

4.3.3 On Economic Inequalities and Bitcoin

While Bitcoin is Inclusive (Sec. 3), i.e., it provides Equal Opportunities, it does not guarantee economic equality. However, economic inequalities, in particular the very glaring ones, are much easier to address with Bitcoin as planetary money than in the current system. We provide several arguments towards supporting this claim:

1. Bitcoin allows transfer of value over the internet in a decentralized, permissionless and inclusive way. Note that this is not possible with the current banking system, as large populations of the world are unbanked and authorities can stop and censor transactions. It is unprecedentedly easy to use this

feature of Bitcoin to “equalize” wealth. A straightforward example, which applies to the world of today, are money remittances, i.e., funds sent by migrant workers (typically in wealthier countries) to friends and family in their homeland. This particular use case has already been a motivating use case in Bitcoin’s adoption in El Salvador, which became the first country in the world to adopt Bitcoin as a legal tender on 7 September 2021.

2. Under the Bitcoin monetary standard, it is very easy for individuals and businesses to understand if they have a disproportionately high fraction of Bitcoin compared to others. They simply need to divide the total supply by their holdings and compare to the number of people on Earth. Note that this is impossible, in general and for the long-term, in the current monetary system regardless of ones’ balance in the bank, as the supply is long-term uncapped (at any given point in time this is possible to estimate, but requires cumbersome computation of the entire fiat money supply and wealth in the world).

   With such an understanding of their own financial security, people could more easily decide to *give and donate their money* thereby reducing economic inequalities.

3. Bitcoin mining can be done practically by anyone, anywhere. While Bitcoin mining consumes energy, energy is very democratically and rather equally distributed across the entire world. In practically every corner of the planet, the sun shines, winds blow, rivers flow and/or geothermal energy exists. Bitcoin’s reliance on energy expenditure therefore incentivizes research on utilizing energy which is available in a respective corner of the planet. Across the planet, renewable resources (as an aggregate) are the most equally distributed ones.

4.4 Conclusions on the Bitcoin Energy Expenditure

In the previous section, we outlined arguments which aim at justifying the energy expenditure of Bitcoin. We painted a reality which, hopefully, provokes the reader to reconsider whether Bitcoin’s energy expenditure is genuinely “useless”, “wasteful” and “excessive”. We did this in an unapologetic way, refraining from even calculating what fraction of world’s energy Bitcoin uses or will use. We maintain that, if Bitcoin brings us closer to a more sustainable world of free people, any fraction of the energy we produce as humans towards this end is justified. A very large energy expenditure on Bitcoin might, perhaps, even help catalyze our evolution to a Type I civilization on the Kardashev scale.[12] The widely-used

[12] https://en.wikipedia.org/wiki/Kardashev_scale

comparison of Bitcoin's energy expenditure to today's nation-states[13] should perhaps only be used to understand which nation-states would still have the power to break the security of Bitcoin (by mounting the 51% attack) if they would somehow engage all the power available in their country.

That said, we should continue investing in research on technologies that would reduce energy expenditure of Bitcoin, yet that would at the same time provide the same or stronger security and decentralization guarantees. So far, such a technology has been elusive. It seems that the second law of thermodynamics might be a challenge here, as Bitcoin reliance on repeated computation of an irreversible hash function, which is the key to its security, results in an irreversible heat production [37].

In a different approach, one might ask whether this expended energy might be used for something "actually useful" and not for simple repetitive, "useless" hashing; such efforts exists, but have so far not been successful. For instance, one proposal made up on-the-fly would be to use this energy towards, say, machine-learning and/or data science. Today, these are sometimes used for good purposes but more often simply to better place ads and "improve" businesses performance in the technologically and humanly inferior monetary fiat system we live with. We have to be very careful in what we define as "useful" and how we define value. What is useful at one level of abstraction and in one value system can be seen as a completely useless activity in another value system of reference, which might actually be considerably better for humanity.

Some readers may be motivated to precisely calculate the net effect of Bitcoin's energy consumption and might feel uneasy to accept the high-level arguments we presented here towards showing that *Bitcoin is good for humankind*. Such a calculation would certainly be interesting to see, yet challenging to perform as Bitcoin's energy expenditure will need to be contrasted to the sum of energy expenditure on activities that Bitcoin renders obsolete or reduces considerably. These include but are not limited to: legacy banking system, data centers powering ad-based spending, and the sum of all activities people undertake under coercion from the current monetary system. It would also need to take into account the human time wastes under the current system and its opportunity cost. Intuitively, we conjecture that the net result will be orders of magnitude in favor of Bitcoin.

[13] See e.g., `https://digiconomist.net/bitcoin-energy-consumption`.

5 Towards Decentralized Cloud Computing

5.1 Challenges

In the previous section, we argued for the Bitcoin network as the backbone for the money of the future. Once we embrace this thought and accept that the future will include a very secure and decentralized, albeit low-throughput consensus network, an interesting question arises: *can we leverage Bitcoin consensus in use cases other than money?*

One possibly interesting use case is decentralized cloud computing and, more generally, the internet. Today's internet is becoming increasingly more centralized, with half of the internet traffic in 2019 coming from only 5 internet companies (well-known social network, cloud and content providers), which is to be contrasted to thousands of autonomous systems (ASs) needed to reach this fraction in 2007 [27].

In the light of this galloping centralization of the internet, powered by the current ad-based economy, it is rightful to ask: *how can we decentralize the internet again?* We clearly need economic models and payment networks to power this, e.g., allowing paying content creators directly, instead through ads. While this can be done on Bitcoin layers 2 and higher [29], as already being implemented by Twitter, the question remains: how do we decentralize the rest of the cloud computing infrastructure, namely computation and storage, while not compromising their security?

First steps have been already made towards using blockchain to support decentralized arbitrary computation and data storage. For instance, Filecoin [6] currently allows decentralized storage of immutable data with rather large capacity (over 10 EB). A number of other projects (e.g., Ethereum [1]), make first steps towards allowing development of general-purpose applications ("smart-contracts") and running them on the blockchain. Even though these are currently used mostly for very limited use cases of debatable value, especially under the prospective Bitcoin monetary standard (such as decentralized finance), the first steps have been made.

Evolving these decentralized systems further entails two main problems: security and scaling. In an attempt to scale their system, currently based on Proof-of-Work, the Ethereum community decided to move to Proof-of-Stake as its (local) backbone, proposing a Proof-of-Stake *beacon chain* [23], citing environmental concerns around Proof-of-Work.

However, this poses issues with security and decentralization. In Section 3, we argued why a network based on Proof-of-Stake cannot be used as a Layer 1 of decentralized systems, due to lack of Inclusiveness. Beyond this fundamental issue, Proof-of-Stake protocols suffer from many attack vectors, including “nothing-at-stake” [20], ”long range” [21] and other attacks. The situation seems unfortunate, as Proof-of-Stake (PoS) and, more generally, identity-based Byzantine fault-tolerant (BFT) protocols [53] remain one of the most promising avenues for improving performance of decentralized systems, in addition to peer-to-peer payment channels [29].

On the other hand, it is unrealistic to expect Bitcoin to include smart-contracts on its main network, nor would this make sense given Bitcoin’s low-throughput. Bitcoin is, rightfully so, conservatively evolving, with major upgrades being deployed rather rarely and never as a hard-fork, i.e., in a backwards-incompatible manner. This is critical to the security of the main network, and it’s the correct approach. When it comes to security, less code is more, as more code practically always introduces more vulnerabilities. For the backbone of the world’s monetary network, introduction of potential vulnerabilities is a clear no-go.

This situation is, fortunately, addressable. In the rest of this section:

- We first discuss, in Section 5.2, how to help secure more scalable consensus protocols, leveraging the Bitcoin blockchain. More specifically, we discuss a promising approach to “saving” Proof-of-Stake blockchains by anchoring (checkpointing) their state and weighted memberhip (stake) into the Bitcoin blockchain, protecting them from long-range and nothing-at-stake attacks. This approach will be possible in a scalable manner only after the Bitcoin Taproot (supported by Bitcoin Core v0.21.1+) upgrade takes effect, at block height 709632, expected in mid-November 2021.

- Then, in Section 5.3, we discuss some of the additional challenges in extending this emerging architecture towards a genuine decentralized cloud.

5.2 Bitcoin as a Backbone for PoS/BFT Protocols

Long-range attacks on Proof-of-Stake [21] rely on the inability of a client or other participant in the system, Alice, who disconnects from the system at time t_1 and reconnects at a later time $t_2 > t_1$, to know that validators (i.e., stakeholders or “miners” in a PoS network) who have been legitimate validators at time t_1 and who leave the system (or transfer their stake) at time t' ($t_2 > t' > t_1$), are not to be trusted anymore. These validators can fork Alice, without her being able to recognize the attack even if she is presented with a “valid” chain fork. We illustrate this attack in Figure 1 for the special case of equal weights among validators (for simplicity). This illustration also demonstrates the related “I still work here” attack [9] in permissioned BFT-based blockchains subject to membership changes (reconfiguration).

Steinhoff et al. recently proposed an approach to deal with the long-range / “I still work here” attack by anchoring BFT/PoS membership into Ethereum’s

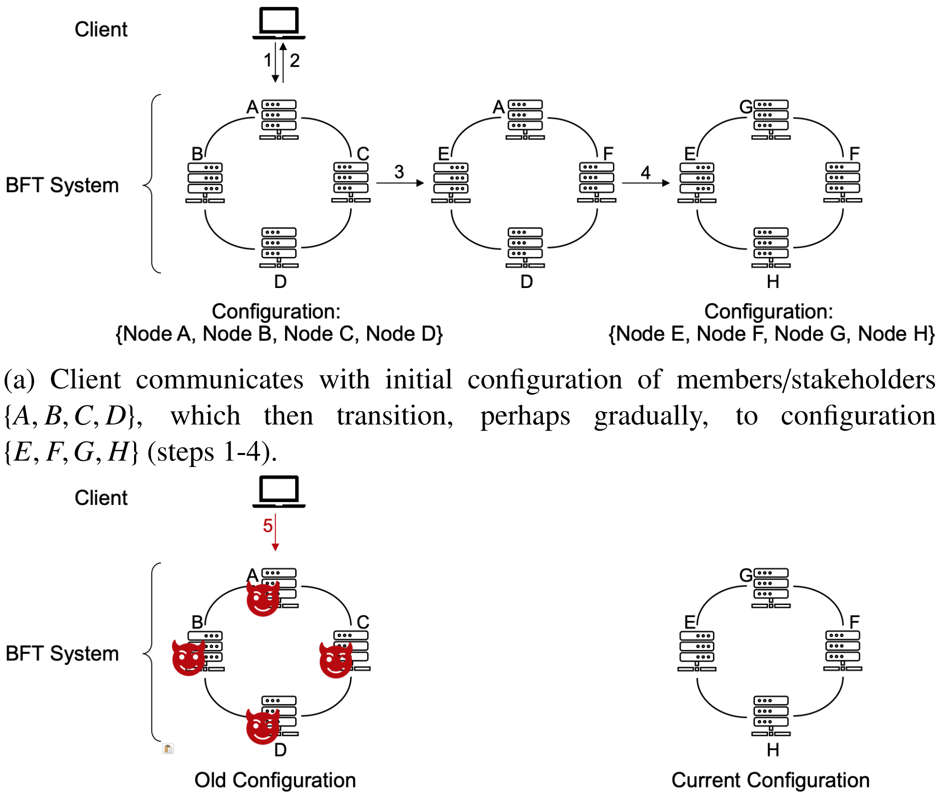

(a) Client communicates with initial configuration of members/stakeholders $\{A, B, C, D\}$, which then transition, perhaps gradually, to configuration $\{E, F, G, H\}$ (steps 1-4).

(b) When client reconnects, the initial configuration may be entirely corrupted by the adversary (step 5).

Figure 1: Illustration of the long-range ("I still work here") family of attacks on PoS (BFT) systems.

Proof-of-Work (PoW) blockchain (Eth 1.0) using Ethereum smart contracts [50]. As a Proof-of-Concept, the approach uses rather unscalable multi-signatures. This approach will no longer be viable once Ethereum abandons Proof-of-Work: the PoS of Eth 2.0 cannot be used instead PoW for anchoring as it too is vulnerable to long-range attacks.

With Bitcoin Taproot[14], one can extend this design and make it scale to thousands of PoS/BFT validators. In short, Bitcoin Taproot's enables Schnorr threshold signatures [45], which permit a public key to be constructed from multiple participant public keys and require cooperation between all participants to sign a transaction. These multi-party keys and signatures are indistinguishable from their single-party equivalents and enable arbitrary t-out-of-n signing policies, very amenable to BFT and PoS protocols.

[14] https://en.bitcoin.it/wiki/BIP_0341

An idea very similar to porting the approach of Steinhoff et al. [50] to Bitcoin Taproot has recently been proposed by Matt Bell [14], along with an initial design, in the context of enabling Proof-of-Stake sidechains for Bitcoin. Indeed, the approach could be used for Bitcoin PoS sidechains which stake in BTC, as well as for "external" PoS blockchains which provide staking in other tokens.

Below, we sketch how this approach might work (different implementations and nuances in the design of this approach are currently being considered — see also [14]):

1. Members of the current configuration of a PoS network A (C_0^A) jointly control, using the Schnorr threshold signature scheme enabled by Bitcoin Taproot, the funds on the Bitcoin blockchain pertaining to network A on address A_0.

2. When the configuration of PoS network A changes sufficiently[15], say to configuration C_1^A, a Bitcoin transaction is constructed, outside the Bitcoin blockchain, as follows:

   (a) Members of the old configuration C_0^A jointly sign a single Bitcoin transaction, which transfers BTC belonging to network A from address A_0 to new address A_1. Here, address A_1 is jointly controlled by members of the new configuration C_1^A.

   (b) The OP_RETURN opcode in the Bitcoin transaction can be used to store a hash of any additional information validators or users of network A need, such as the Merkle root of the most recent weighted stakeholders or even the Merkle root of the entire state of network A.[16]

   (c) Whatever information is referred to by its hash and stored in the OP_RETURN opcode of the Bitcoin transaction, could then be stored in its entirety into an external content-addressable decentralized storage. This could be for example, Filecoin decentralized storage, or the Interplanetary File System (IPFS) [15, 7].

   In this case, the Filecoin/IPFS content ID (CID)[17] of new membership C_1^A, and the latest state checkpoint of network A, could be stored in

[15] Here, the notion of a "sufficient membership change" is intentionally underspecified and may depend on the security policy of network A.

[16] Per https://en.bitcoin.it/wiki/OP_RETURN, "Many members of the Bitcoin community believe that the use of OP_RETURN is irresponsible... for arbitrary data". Fortunately, Bitcoin is inclusive and decentralized and one does not need to ask permission from other members of the community for using Bitcoin in a certain way.

[17] See https://docs.ipfs.io/concepts/content-addressing/.

the Bitcoin transaction OP_RETURN opcode, whereas the data corresponding to said CID could be stored in Filecoin/IPFS.

3. The "account" of network A on the Bitcoin blockchain needs to be periodically refunded with BTC to account for transaction fees on the Bitcoin network. This could be done by having new nodes joining or staking on network A deposit funds to the latest address used by network A on the Bitcoin blockchain (above, address A_1).

4. With such an approach to reconfiguration in place, Alice, the user of PoS network A from the beginning of our long-range attack example, can resolve, by running a Bitcoin full-node and a client of content-addressable decentralized storage (e.g., Filecoin/IPFS), all the information pertaining to PoS network A.

   It is worth noting that running a Bitcoin full node is not an issue here for Alice, due to Bitcoin's high operational decentralization (Sec. 4.2), as a Bitcoin node could be run by an IoT device (e.g., Raspberry Pi) and a standard hard disk. In this context, the fact that Bitcoin blockchain size grows "slowly" due to the small block size and 10-minute block time is an important advantage.

Clearly, in the design outlined above, in addition to using Bitcoin as a stake (membership) anchor, PoS network A could use its own content addressable storage implementation instead of Filecoin/IPFS. However, such implementations are rather involved and it is considerably easier for developers of new PoS networks to leverage existing building blocks for the decentralized internet, i.e., Bitcoin and Filecoin/IPFS.

5.3 Next Steps for Decentralized Computation and Storage

In order for decentralized cloud computing and the decentralized internet (so-called Web3) to become pervasive and supersede current Web2 internet and centralized cloud computing, it is reasonable to require Web3 to run Web2 workloads, with billions of transactions per second, low latencies, high security, etc. All this should be ideally done with higher privacy, censorship resilience (freedom of information) and availability, to provide benefits for humanity over Web2.

This tasks is very challenging, in particular in the context of scalability. However, we firmly believe it is achievable. Towards this goal, we highlight some of the decentralized systems areas which require more research. In a non-exhaustive approach, focusing on problems related to scaling distributed computing for Web3, we divided the problems to be tackled into the following research areas: 1) hierarchical scaling, 2) scaling consensus proper, and 3) scalable execution. These are supported by research considerations that are pervasive to all of these work areas: security (including privacy), decentralization, and correctness of design and implementation.

Hierarchical scaling. Assuming a future Web3 handling Web2-sized workloads, one cannot rely on a single L2 blockchain network (assuming Bitcoin at L1) to “rule them all”, much like today’s web workloads are not executed on any single centralised machine.

This immediately brings sharding to the picture, i.e. the horizontal scaling of decentralized systems. This requires extending the work at the intersection of classical distributed computing and database problems related to concurrency — such as cross-shard transaction semantics and shard state management — to Web3-specific security challenges. There is already considerable work in this area (see, e.g., [38, 32, 56] and the surveys [57, 55]), which is picking up in pace in recent years, with more interesting and creative proposals appearing recently (e.g., [34, 39]). Sharding challenges related to scalability and security of shards are complemented by low-level interoperability [13] challenges among consensus protocols of different families.

Not precluding alternative designs, we envision, as an extension of our designed outlined in Section 5.2, a world of *hierarchical consensus protocols* in which children shards, powered by faster and more performant consensus protocols, leverage stronger security of their parent shards, which might in turn have weaker performance. In such a hierarchical approach, checkpointing into a higher security context will be important.

Scaling consensus proper. In a possible hierarchical consensus approach we mentioned above, different consensus protocols will be applicable to different shards, layers and use cases.

We propose identifying the “best” consensus protocol within a given security and scalability context, applicable to a single shard. Different consensus protocols may be relevant, e.g., targeting distinct sybil attack protection mechanisms (e.g., based on verifiable resource commitments [42, 41] or Proof-of-Stake), participating population size (e.g., whether we are dealing with thousands [28] or hundreds [49] of nodes), performance, and security guarantees. Here, we should pay particular attention to the tradeoffs between ease of design, maintainability, and system guarantees.

Here, it is very important to understand that scalable consensus protocol implementations need to evolve into full-fledged battle-tested production systems. There is work ahead of us not only in protocol design but in robust testing and verification tools.

Scaling execution. Advances in scaling through sharding and scaling consensus are insufficient without considering the scalability of application execution. Existing blockchain systems largely follow, within a single shard, a classical order-execute architecture in which sequential execution of applications (e.g. smart contracts or payment scripts) follows prior ordering of transactions [54]. This introduces severe performance bottlenecks, which have been the target of decades of research in databases, multi-core processors, and distributed systems.

Here, we will need to revisit the parallel execution problem through the prism of existing decentralised applications (based on dedicated smart-contract virtual machines [22]) but also aim at accommodating larger-scale general-purpose scale computations, including federated machine learning workloads or computations over large data [33, 16], e.g., stored on IPFS. This may involve porting alternative execution models, such as execute-order-validate, tested in the domain of permissioned blockchains [11] to the world of permissionless systems. Alternative application programming models which support better efficiency of parallel execution (e.g., CRDTs [46]) are also of particular interest.

Finally, very relevant to all these research areas in particular hierarchical scaling and consensus proper, will be further research on improving distributed randomness used in actual decentralized systems (e.g., drand [5]), using perhaps Bitcoin itself [18], or approaches based, e.g., on verifiable delay functions [17].

6 Conclusions and Future Work

In this paper we argued for Bitcoin as a backbone of future decentralized money and the internet, making a case for it to be good for humankind.

The main novelty of this paper is not necessarily technical — the main contribution is a proposal to give existing systems, notably Bitcoin, a different look. Author's own view on Bitcoin changed only relatively recently from "generally positive" to the one depicted in this paper, and this was equal to an enlightenment. Arguably, one cannot change other people, but can only change oneself. The goal of this paper is simply to share the message.

Future work in building decentralized systems, briefly tackled in Section 5.3, will, by definition, not be a task of any single individual or organization. An interesting question here is how do we, decentralized systems researchers and developers, organize better towards building coherent and interconnected systems, while remaining in a decentralized working mode and outside the control of any single organization?

In an attempt to facilitate this, we recently established ConsensusLab as a research and development hub and part of Protocol Labs. One of the main goals of ConsensusLab is to serve as a weak synchronization point for researchers across academia, open-source blockchain projects and industry, and to help foster discussion and open collaboration pertaining to decentralized consensus and related technologies that are at the heart of decentralized systems. The collaboration here is, in part, related to openly discussing, criticizing and, perhaps, gently steering technical innovations in consensus, but is also very much related to meta-collaboration, discussing how do we build tools, incentives and funding schemes towards working together in a decentralized way. We wholeheartedly look forward to these collaborations and future innovations.

Acknowledgments

I am indebted to Seth Gilbert and Stefan Schmid, editors of the BEATCS Distributed Computing column, for providing me with an invaluable opportunity, inviting my contribution, catalyzing the writeup of this position paper and giving valuable feedback. I thank Evan Miyazono for very useful comments, including questioning my arguments. I am deeply grateful to Jorge Soares, whose very detailed feedback helped considerably improve the paper. I thank Juan Benet for inspiring discussions around hierarchical consensus and for proposing ambitious goals for Web3.

Finally, I wish to express my deep gratitude to the people who helped me come to the viewpoint expressed in this paper. I am indebted to Andreja Čobeljić for helping me better understand the world around me and many discussions on cryptocurrencies and beyond. I thank Bojana and Nenad Mijailović for referring me to the works of the economists of the Austrian School. Finally, I heartily thank Ana Vukolić for endless discussions and invaluable feedback on the argument presented here.

References

[1] Ethereum. http://ethereum.org.

[2] Proof of stake instead of proof of work. https://bitcointalk.org/index.php?topic=27787.0, 2011.

[3] I AM HODLING. https://bitcointalk.org/index.php?topic=375643.0, 2013.

[4] Cambridge Bitcoin electricity consumption index. https://cbeci.org/cbeci/comparisons, 2021.

[5] drand: Distributed randomness beacon. https://drand.love/, 2021.

[6] Filecoin. https://filecoin.io/, 2021.

[7] Interplanetary file system (ipfs). https://ipfs.io/, 2021.

[8] Wikipedia: Richard Cantillon. https://en.wikipedia.org/wiki/Richard_Cantillon, 2021.

[9] M. K. Aguilera, I. Keidar, D. Malkhi, J. Martin, and A. Shraer. Reconfiguring replicated atomic storage: A tutorial. *Bull. EATCS*, 102:84–108, 2010.

[10] B. Alpern and F. B. Schneider. Recognizing safety and liveness. *Distributed Comput.*, 2(3):117–126, 1987.

[11] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. D. Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolic, S. W. Cocco, and J. Yellick. Hyperledger fabric: a distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference, EuroSys 2018, Porto, Portugal, April 23-26, 2018*, pages 30:1–30:15, 2018.

[12] S. Azouvi, M. Maller, and S. Meiklejohn. Egalitarian society or benevolent dictatorship: The state of cryptocurrency governance. In A. Zohar, I. Eyal, V. Teague, J. Clark, A. Bracciali, F. Pintore, and M. Sala, editors, *Financial Cryptography and Data Security - FC 2018 International Workshops, BITCOIN, VOTING, and WTSC, Nieuwpoort, Curaçao, March 2, 2018, Revised Selected Papers*, volume 10958 of *Lecture Notes in Computer Science*, pages 127–143. Springer, 2018.

[13] R. Belchior, A. Vasconcelos, S. Guerreiro, and M. Correia. A survey on blockchain interoperability: Past, present, and future trends, 2021.

[14] M. Bell. Proof-of-stake bitcoin sidechains. https://gist.github.com/mappum/da11e37f4e90891642a52621594d03f6, 2021.

[15] J. Benet. IPFS - content addressed, versioned, P2P file system. *CoRR*, abs/1407.3561, 2014.

[16] K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konečný, S. Mazzocchi, H. B. McMahan, T. V. Overveldt, D. Petrou, D. Ramage, and J. Roselander. Towards federated learning at scale: System design, 2019.

[17] D. Boneh, J. Bonneau, B. Bünz, and B. Fisch. Verifiable delay functions. In *Advances in Cryptology – CRYPTO 2018*, volume 10991 of *Lecture Notes in Computer Science*, pages 757–788. Springer, 2018.

[18] J. Bonneau, J. Clark, and S. Goldfeder. On Bitcoin as a public randomness source. *IACR Cryptology ePrint Archive*, 2015:1015, 2015.

[19] N. Bostrom. Are we living in a computer simulation? *Philosophical Quarterly*, 53(211), 2003.

[20] J. Brown-Cohen, A. Narayanan, A. Psomas, and S. M. Weinberg. Formal barriers to longest-chain proof-of-stake protocols. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, page 459–473, New York, NY, USA, 2019. Association for Computing Machinery.

[21] E. Deirmentzoglou, G. Papakyriakopoulos, and C. Patsakis. A survey on long-range attacks for proof of stake protocols. *IEEE Access*, 7:28712–28725, 2019.

[22] T. D. Dickerson, P. Gazzillo, M. Herlihy, and E. Koskinen. Adding concurrency to smart contracts. *Distributed Comput.*, 33(3-4):209–225, 2020.

[23] Ethereum Foundation. Proof-of-stake (PoS). https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/, July 2021.

[24] G. C. Fanti, L. Kogan, S. Oh, K. Ruan, P. Viswanath, and G. Wang. Compounding of wealth in proof-of-stake cryptocurrencies. In I. Goldberg and T. Moore, editors, *Financial Cryptography and Data Security - 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18-22, 2019, Revised Selected Papers*, volume 11598 of *Lecture Notes in Computer Science*, pages 42–61. Springer, 2019.

[25] P. D. Filippi and B. Loveluck. The invisible politics of Bitcoin: Governance crisis of a decentralized infrastructure. *Internet Policy Review*, 5(4), 2016.

[26] A. E. Gencer, S. Basu, I. Eyal, R. van Renesse, and E. G. Sirer. Decentralization in Bitcoin and Ethereum networks. In S. Meiklejohn and K. Sako, editors, *Financial Cryptography and Data Security - 22nd International Conference, FC 2018, Nieuwpoort, Curaçao, February 26 - March 2, 2018, Revised Selected Papers*, volume 10957 of *Lecture Notes in Computer Science*, pages 439–457. Springer, 2018.

[27] P. Gigis, M. Calder, L. Manassakis, G. Nomikos, V. Kotronis, X. Dimitropoulos, E. Katz-Bassett, and G. Smaragdakis. Seven years in the life of hypergiants' off-nets. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, SIGCOMM '21, page 516–533, New York, NY, USA, 2021. Association for Computing Machinery.

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

[29] L. Gudgeon, P. Moreno-Sanchez, S. Roos, P. McCorry, and A. Gervais. Sok: Layer-two blockchain protocols. In J. Bonneau and N. Heninger, editors, *Financial Cryptography and Data Security - 24th International Conference, FC 2020, Kota Kinabalu, Malaysia, February 10-14, 2020 Revised Selected Papers*, volume 12059 of *Lecture Notes in Computer Science*, pages 201–226. Springer, 2020.

[30] F. A. Hayek. *The Constitution of Liberty: The Definitive Edition*. The University of Chicago Press, 1960.

[31] D. Karakostas, A. Kiayias, C. Nasikas, and D. Zidros. Cryptocurrency egalitarianism:a quantitative approach. In *International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2019)*, 2019.

[32] E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford. Omniledger: A secure, scale-out, decentralized ledger via sharding. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*, pages 583–598. IEEE Computer Society, 2018.

[33] J. Konečny, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon. Federated learning: Strategies for improving communication efficiency, 2017.

[34] M. Król, O. Ascigil, S. Rene, A. Sonnino, M. Al-Bassam, and E. Rivière. Shard scheduler: object placement and migration in sharded account-based blockchains, 2021.

[35] L. Lamport. Proving the correctness of multiprocess programs. *IEEE Trans. Software Eng.*, 3(2):125–143, 1977.

[36] L. Lamport, R. E. Shostak, and M. C. Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.

[37] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[38] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena. A secure sharding protocol for open blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, page 17–30, New York, NY, USA, 2016. Association for Computing Machinery.

[39] G. A. Marson, S. Andreina, L. Alluminio, K. Munichev, and G. Karame. Mitosis: Practically scaling permissioned blockchains, 2021.

[40] A. Miller. *Permissioned and Permissionless Blockchains*, pages 193–204. 2019.

[41] T. Moran and I. Orlov. Simple proofs of space-time and rational proofs of storage. Cryptology ePrint Archive, Report 2016/035, 2016. https://ia.cr/2016/035.

[42] Protocol Labs. Filecoin: a decentralized storage network. https://filecoin.io/filecoin.pdf, 2017.

[43] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. https://bitcoin.org/bitcoin.pdf, 2008.

[44] A. R. Sai, J. Buckley, B. Fitzgerald, and A. LeGear. Taxonomy of centralization in public blockchain systems: A systematic literature review. *Information Processing and Management*, 58(4), July 2021.

[45] C. P. Schnorr. Efficient identification and signatures for smart cards. In G. Brassard, editor, *Advances in Cryptology — CRYPTO' 89 Proceedings*, pages 239–252, New York, NY, 1990. Springer New York.

[46] M. Shapiro, N. M. Preguiça, C. Baquero, and M. Zawirski. Conflict-free replicated data types. In X. Défago, F. Petit, and V. Villain, editors, *Stabilization, Safety, and Security of Distributed Systems - 13th International Symposium, SSS 2011, Grenoble, France, October 10-12, 2011. Proceedings*, volume 6976 of *Lecture Notes in Computer Science*, pages 386–400. Springer, 2011.

[47] H. C. Simons. Rules versus authorities in monetary policy. *Journal of Political Economy*, 44:1–30, 1936.

[48] B. Srinivasan. Quantifying decentralization. Blockstack Summit 2017, https://www.youtube.com/watch?v=4UXT5YVJwB4, 2017.

[49] C. Stathakopoulou, T. David, M. Pavlovic, and M. Vukolić. Mir-BFT: High-throughput robust BFT for decentralized networks, 2021.

[50] S. Steinhoff, C. Stathakopoulou, M. Pavlovic, and M. Vukolić. BMS: Secure Decentralized Reconfiguration for Blockchain and BFT Systems, 2021.

[51] A. S. Tanenbaum and M. van Steen. *Distributed systems - principles and paradigms, 2nd Edition*. Pearson Education, 2007.

[52] C. Troncoso, M. Isaakidis, G. Danezis, and H. Halpin. Systematizing decentralization and privacy: Lessons from 15 years of research and deployments. *Proc. Priv. Enhancing Technol.*, 2017(4):404–426, 2017.

[53] M. Vukolić. The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication. In *International Workshop on Open Problems in Network Security (iNetSec)*, pages 112–125, 2015.

[54] M. Vukolić. Rethinking permissioned blockchains. In *Proceedings of the ACM Workshop on Blockchain, Cryptocurrencies and Contracts*, BCC '17, page 3–7, New York, NY, USA, 2017. Association for Computing Machinery.

[55] G. Wang, Z. J. Shi, M. Nixon, and S. Han. Sok: Sharding on blockchain. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, AFT '19, page 41–61, New York, NY, USA, 2019. Association for Computing Machinery.

[56] J. Wang and H. Wang. Monoxide: Scale out blockchains with asynchronous consensus zones. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, pages 95–112, Boston, MA, Feb. 2019. USENIX Association.

[57] G. Yu, X. Wang, K. Yu, W. Ni, J. A. Zhang, and R. P. Liu. Survey: Sharding in blockchains. *IEEE Access*, 8:14155–14181, 2020.