

The Quest for Scalable Blockchain Fabric: Proof-of-Work vs. BFT Replication

Marko Vukolić^(✉)

IBM Research, Zurich, Switzerland
mvu@zurich.ibm.com

Abstract. Bitcoin cryptocurrency demonstrated the utility of global consensus across thousands of nodes, changing the world of digital transactions forever. In the early days of Bitcoin, the performance of its probabilistic *proof-of-work* (PoW) based consensus fabric, also known as *blockchain*, was not a major issue. Bitcoin became a success story, despite its consensus latencies on the order of an hour and the theoretical peak throughput of only up to 7 transactions per second.

The situation today is radically different and the poor performance scalability of early PoW blockchains no longer makes sense. Specifically, the trend of modern cryptocurrency platforms, such as Ethereum, is to support execution of arbitrary distributed applications on blockchain fabric, needing much better performance. This approach, however, makes cryptocurrency platforms step away from their original purpose and enter the domain of database-replication protocols, notably, the classical *state-machine replication*, and in particular its Byzantine fault-tolerant (BFT) variants.

In this paper, we contrast PoW-based blockchains to those based on BFT state machine replication, focusing on their scalability limits. We also discuss recent proposals to overcoming these scalability limits and outline key outstanding open problems in the quest for the “ultimate” blockchain fabric(s).

Keywords: Bitcoin · Blockchain · Byzantine fault tolerance · Consensus · Proof-of-work · Scalability · State machine replication

1 Introduction

Distributed consensus, infamous for its limited scalability, was for decades perceived as a synchronization primitive that is to be used only in applications in desperate need of consistency and only among few nodes (see e.g., [8,27]). However, Nakamoto’s Bitcoin cryptocurrency [47] demonstrated the utility of decentralized consensus across thousands of nodes, changing the world of digital transactions forever.

Although the Bitcoin protocol does not actually implement consensus in the traditional distributed computing sense, it comes very close to consensus with

probabilistic agreement [25]. In a nutshell, the goal of a cryptocurrency such as Bitcoin, is to totally order transactions on a distributed ledger, also called a *blockchain*. The Bitcoin blockchain consists of a hashchain of blocks: every block contains an ordered set of transactions and a hash of the preceding block (starting from the initial, the so-called “genesis” block). The key part is the *Proof-of-Work* (PoW) aspect of the hashchain [21]: a Bitcoin block contains nonces that a Bitcoin *miner* (i.e., a node attempting to add a block to the chain) must set in such a way that the hash of the entire block is smaller than a known *target*, which is typically a very small number. In fact, in Bitcoin, the *difficulty* of mining, inversely proportional to the target, is adjusted dynamically throughout the lifetime of the system. The adjustment is made with respect to the block-mining rate and, indirectly, with respect to the computational power of nodes participating in the system, to maintain the expected block-mining rate at roughly one block every 10 min [47]. This latency of 10 min (per block) is often referred to as the *block frequency* (see e.g., [22]) and is one of the two critical “magic numbers” in Bitcoin, the other being the *block size*, which is set in Bitcoin to 1 MB.

In the early days of Bitcoin, the *performance scalability* of its probabilistic PoW-based blockchain was not a major issue. Even today, Bitcoin works with a consensus latency of about an hour (for the recommended 6-block transaction confirmation), and with up to 7 (seven) transactions per second peak throughput (with smallest 200–250 byte transactions). On top of this, the Bitcoin network uses a lot of power, which, in 2014, was roughly estimated to be in the ballpark of 0.1–10 GW [48].

However, blockchain requirements change rapidly, with high latency and low throughput of Bitcoin-like blockchain becoming a major challenge [6]. As a comparison, leading global credit-card payment companies serve roughly 2000 transactions per second on average [58], with a peak capacity designed to sustain more than 10000 transactions per second. Moreover, the trend of modern cryptocurrency platforms, such as Ethereum [57], is to support execution of Turing-complete code on blockchain fabric in the form of *smart contracts*, which are, roughly speaking, custom, self-executing programs (distributed applications) that automatically enforce properties of a digital contract. In fact, smart-contract blockchain is seen as a candidate technology for distributed ledgers in many industries. Clearly, in many of the intended smart-contract use cases, distributed applications require much better performance than that offered by Bitcoin. The banking industry is one prominent example, where potential blockchain use cases go well beyond digital payments [45] to, e.g., securities trade settlements and trade finance.

Smart-contract use cases take the blockchain well beyond its original cryptocurrency purpose, back to the domain of database replication protocols, notably, the classical *state-machine replication* [53]. Indeed, a smart contract can be modeled as a state machine, and its consistent execution across multiple nodes in a distributed environment can be achieved using state machine replication. A family of state-machine replication protocols particularly interesting

for blockchain is the family of *Byzantine fault-tolerant* (BFT) [37] state-machine replication protocols, which promise consensus despite participation of malicious (Byzantine) nodes. In more than three decades of research, BFT protocol prototypes have been shown to be practical [10], reaching practically minimal latencies allowed by the network, and supporting tens of thousands transactions per second (see e.g., [3,34]). However, BFT and state-machine replication protocols in general are often challenged for their scalability in terms of number of nodes (replicas) [8], and have not been thoroughly tested in this aspect critical to blockchain.

In summary, blockchain consensus technologies of today, PoW and BFT, sit at the two opposite ends of the scalability spectrum. Roughly speaking, PoW-based blockchain offers good node scalability with poor performance, whereas BFT-based blockchain offers good performance for small numbers of replicas, with not-well explored and intuitively very limited scalability. This current state of blockchain scalability is sketched in Fig. 1. Given seemingly inherent tradeoffs between the number of replicas and performance, it is not clear today what the optimal blockchain solution is for the sweet spot relevant for many use cases in which the number of nodes n ranges from a few tens to 1000 (or perhaps few thousands).

In this paper, we overview recent efforts towards improving scalability on both sides of the spectrum and highlight interesting directions and open problems in the quest for the “ultimate” blockchain fabric. First, in Sect. 2 we com-




Fig. 1. Illustration of performance and scalability of different families of PoW and BFT protocols discussed in this paper. The actual, real-world performance of systems that touch upon the grey area is subject to further research. Hence, their positioning within the grey area is at the moment entirely speculative and for motivational purposes only.

pare PoW-based blockchains to those based on BFT state-machine replication. Then, in Sect. 3, we overview novel promising approaches to scaling PoW and BFT protocols. We conclude in Sect. 4 with several open questions that will be interesting to tackle in the very near future.

2 PoW vs. BFT Blockchains

Table 1 gives a high-level comparison between PoW consensus and BFT consensus for a set of important blockchain properties. These properties include node identity management, consensus finality (or, dually, the possibility of temporary forks in the blockchain), scalability in terms of number of consensus nodes and clients, performance (latency, throughput, power consumption), tolerated power of adversary, network synchrony assumptions, and, last but not least, existence of correctness proofs of protocols underlying blockchain. This set of properties is certainly not exhaustive, but we believe it is representative for comparing two blockchain families. In the rest of this section, we discuss Table 1 in more detail.

Node Identity Management. How node identities are managed in PoW and BFT protocols is possibly their most fundamental difference. PoW blockchains fea-

Table 1. High-level comparison between PoW and BFT blockchain consensus families for a set of important blockchain properties. Entries in bold suggest desirable features and highlight advantages of one consensus family over the other.

	PoW consensus	BFT consensus
Node identity management	open, entirely decentralized	permissioned, nodes need to know IDs of all other nodes
Consensus finality	no	yes
Scalability (no. of nodes)	excellent (thousands of nodes)	limited, not well explored (tested only up to $n \leq 20$ nodes)
Scalability (no. of clients)	excellent (thousands of clients)	excellent (thousands of clients)
Performance (throughput)	limited (due to possible of chain forks)	excellent (tens of thousands tx/sec)
Performance (latency)	high latency (due to multi-block confirmations)	excellent (matches network latency)
Power consumption	very poor (PoW wastes energy)	good
Tolerated power of an adversary	$\leq 25\%$ computing power	$\leq 33\%$ voting power
Network synchrony assumptions	physical clock timestamps (e.g., for block validity)	none for consensus safety (synchrony needed for liveness)
Correctness proofs	no	yes

ture an entirely decentralized identity management — for example, anybody can download the code for Bitcoin miner, and start participating in the protocol, knowing basically only a single peer to start with. This is a very powerful feature of PoW blockchains and the main reason why they are the blockchain family of choice when it comes to so-called “public” blockchains in which anybody is allowed to participate. Such public blockchains are sometimes also called “permissionless” blockchains — permissionless participation is made possible by PoW, as PoW inherently addresses the Sybil attack [18], infamous in anonymous networks. Specifically, in PoW-based blockchains, the ability of a node (resp., a pool of nodes) to influence the outcome of PoW consensus depends on computational power of a node (resp., a pool).

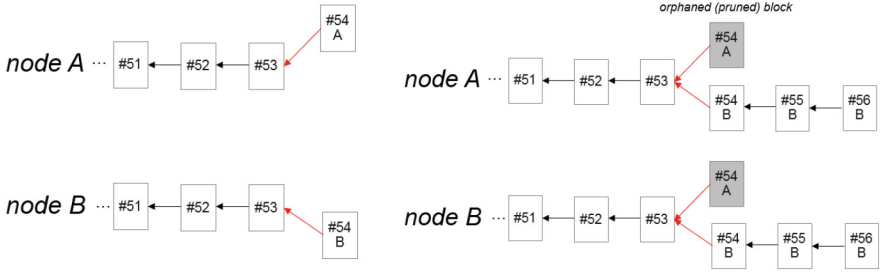
In contrast, the BFT approach to consensus typically requires every node to know the entire set of its peer nodes participating in consensus. This in turn calls for a (logically) centralized identity management in which a trusted party issues identities and cryptographic certificates to nodes.¹ Intuitively, this aspect of BFT-based blockchains puts it at a disadvantage with respect to PoW blockchains. That said, in a number of emerging blockchain applications (e.g., banking, finance, land and real-estate ownership ledgers) the requirement for known identity of nodes might anyway be imposed for legal and compliance reasons. This explains why BFT consensus protocols are the technology of choice for so-called “permissioned” blockchains, which require blockchain participants identity to be known.

Consensus Finality. Roughly speaking, what is often informally referred to as “consensus finality” (and sometimes as “forward security” [15]) is a property that mandates that a valid block, appended to the blockchain at some point in time, be never removed from the blockchain. In the standard distributed computing terminology, “consensus finality” follows from a combination of the *total order* and *agreement* properties of total order (atomic) broadcast [17], which is the primitive all state-machine replication protocols are built upon (total order broadcast is, in turn, equivalent to consensus). Translated to blockchain terminology, this property can be phrased as follows:

Definition 1 (Consensus Finality). *If a correct node p appends block b to its copy of the blockchain before appending block b' , then no correct node q appends block b' before b to its copy of the blockchain.*

Consensus finality is not satisfied by PoW-based blockchains. To see why, note that, besides obviating the need for identity management, PoW acts as a randomized concurrency control mechanism, in which the block frequency is adjusted such that block collisions (i.e., concurrent appends of different blocks to the blockchain) are rare. However, as concurrency control is only probabilistic and as block propagation over a network can take some time [16], collisions

¹ Here, it is important to note that after an initial bootstrap of a BFT-based blockchain, the nodes already on the blockchain could themselves act together as a distributed trusted party and help reconfigure the system [5, 52].



(a) Consensus finality violation resulting in a fork. (b) Eventually, one of the blocks must be pruned by a conflict resolution rule (e.g., Bitcoin's longest chain rule).

Fig. 2. Illustration of a violation of consensus finality, fork and conflict resolution.

do happen, resulting in temporary forks on the blockchain that PoW-based blockchains are prone to even if all nodes are honest. These temporary forks (see Fig. 2 for an illustration) are resolved by rules such as Bitcoin's longest (most difficult) fork rule [47], or the GHOST rule [54], a variant of which is used in Ethereum. However, the very presence of temporary forks implies no consensus finality. As we discuss in more detail below, absence of consensus finality directly impacts the consensus latency of PoW blockchains as transactions need to be followed by several blocks to increase the probability that a transaction will not end up being pruned and removed from the blockchain (we speak of multi-block confirmation).

In contrast, consensus finality is satisfied by all BFT and state-machine replication protocols.² This gives BFT-based blockchains a clear advantage over PoW, as applications, users and smart contracts can have immediate confirmation of the final inclusion of a transaction into the blockchain.

Scalability. Although decoupling the issue of blockchain scalability (with the number of nodes and clients in the system) from that of blockchain performance (latency and throughput) is not entirely possible, we nevertheless first focus on the number of nodes and clients for which PoW and BFT technologies have been proven to work in practice.

On the one hand, the Bitcoin network features thousands of mining nodes, demonstrating node scalability of PoW-based blockchains in practice. That said, it is worth mentioning that grouping of miners into mining pools (with the goal of splitting mining rewards and making mining a financially more predictable endeavour) plagues Bitcoin, effectively centralizing the cryptocurrency [26]. We note that mining pool centralization is not a specific trait of Bitcoin, but more a consequence of the popularity of a PoW blockchain, affecting also many *altcoins* (alternative Bitcoin-like cryptocurrencies) as well as popular blockchains, such as Ethereum.

² Provided the assumptions about the power of the adversary hold.

On the other hand, BFT and state-machine replication are, in general, perceived as protocols with poor scalability (see, e.g., Brewer’s CAP theorem [8]). However, having been invented in the context of replicating traditional applications, such as databases, for fault-tolerance, BFT protocols were never really tested thoroughly for their scalability beyond, say, $n = 10$ or $n = 20$ nodes, in particular in the light of the fairly modest performance targets of many blockchain applications. Intuitively, because of their intensive network communication which often involves as many as $O(n^2)$ messages per block [10], BFT protocols are seen in the database and systems communities as not scalable (see also [44]).³ This is true even for their crash-tolerant counterparts, i.e., replication protocols such as Paxos [36], Zab [30] and Raft [49], which are used in many large scale systems but practically never across more than a handful of replicas (see e.g., [13]).

Finally, when it comes to scalability with the number of clients, both PoW and BFT protocols support thousands of clients and scale well.

Performance. Beyond the very limited performance of Bitcoin of up to 7 transactions per second (with the current block size) and 1 h latency with 6-block confirmation, PoW-based blockchains face inherent performance challenges. As we already discussed, the two main performance-related parameters of a PoW blockchain are *block size* and *block frequency*. Increasing the block size with the goal of boosting throughput comes at the cost of increasing the latency, because of longer propagation delays of larger blocks across the Internet. These longer delays, in turn, have negative implications on blockchain security: longer delays may increase the number of forks and the possibilities for mounting double-spending attacks [33], because of the possibility of temporary chain forks and absence of consensus finality in PoW blockchains. Similar security challenges apply when the block frequency is increased, with the goal of reducing the latency of multi-block confirmation. The exact security implications of tuning the block frequency and the block size in PoW-based blockchain are in general rather involved (see e.g., [54] for an analysis) and should be handled with care. With this in mind, limited performance is seemingly inherent to PoW blockchains and not an artifact of a particular implementation.

In contrast, modern BFT protocols have been confirmed to sustain tens of thousands of transactions with practically network-speed latencies, not only as prototypes (e.g., [3, 12, 34]) but also as practical systems [5].

Adversary. PoW and BFT consider different adversaries. In PoW blockchains, what matters is the total computational (hashing) power controlled by the adversary. Initially, Bitcoin was thought to be invulnerable so long as the adversary controls less than 50% of hashing power. Years later, it was shown that Bitcoin mining is actually vulnerable even if only 25% of the computing power is

³ That said, it is worth noting that there are *optimistic* BFT protocols with $O(n)$ common-case (expected) message-complexity (see, e.g., [3, 51]) — we discuss these later in more detail.

controlled by an adversary [23]. In contrast, BFT voting schemes are known to tolerate at most $n/3$ corrupted nodes [20]. This bound holds only when the network is allowed to be (from time to time) fully asynchronous — strengthening synchrony assumptions makes it possible to raise this threshold. The classical $n/3$ threshold bound for BFT consensus can be generalized to general adversary structures, where an adversary can control different subsets of nodes [28, 56].

Network Synchrony. Bitcoin relies on the local time of a node to timestamp a block. Roughly speaking, a block is accepted as valid if its timestamp is greater than the median of the last 11 blocks. Additionally, timestamps play a major role in calculating the difficulty of mining and maintaining block frequency. Therefore, loose clock synchrony is needed for liveness. However, timestamp manipulation attacks that may also compromise the consistency of the blockchain are conceivable (see the “zeitgeist attack” [1]). Although such attacks are difficult to stage against major PoW blockchains such as Bitcoin, they have been successfully performed in the context of some PoW altcoins.

BFT protocols typically do not rely on any physical clock.⁴ However, eventually synchronous communication is needed to ensure liveness, owing to the FLP consensus impossibility result, which states that consensus is impossible to achieve deterministically with potentially faulty nodes in a purely asynchronous system [24]. The safety properties of consensus, including consensus finality, are maintained despite global communication outages and arbitrarily long asynchrony periods [20].

Correctness Proofs. Historically, state-machine replication protocols, and in particular their BFT variants, have been recognized as very challenging to design and implement [3, 5, 11]. Consequently, new protocols are subject to detailed academic scrutiny and therefore come with (more or less) detailed proofs, sometimes even with formal proofs that take an entire PhD thesis (see [14, 40]). Even if it may be understandable why Bitcoin was originally deployed without having been subjected to similar scrutiny, it is rather surprising that novel PoW blockchains are rarely accompanied by a detailed security and distributed protocol and security analysis.

3 Improving Blockchain Scalability

In this section we overview and discuss several recent efforts that focus on improving the scalability aspects of both PoW and BFT blockchains.

Improving the Performance of PoW Blockchains. Sompolinski and Zohar recently proposed the GHOST (Greedy Heaviest-Observed Sub-Tree) rule [54], which basically resolves conflicts in a PoW blockchain by weighing the subtrees rooted in blocks rather than the longest (sub)chain rooted in given blocks.

⁴ Some state-machine replication protocols do use physical clock timestamps, but only to improve performance [19].

Although GHOST is essentially a conflict-resolution strategy, it offers performance benefits over the standard longest (heaviest) chain rule of Bitcoin, as it provides more secure means of increasing the block frequency and the block size [54]. A variant of the GHOST rule is actually implemented in the Ethereum blockchain [57], although the GHOST-PoW performance has not yet been adequately stress-tested with high loads (in 2016, typical Ethereum throughput is fewer than 20,000 transactions per day, i.e., about 0.2 tx/s on average).⁵

Bitcoin-NG is a novel proposal by Eyal et al. [22] that uses standard PoW for leader election, declaring a node which mines a block with standard difficulty (called a key block) to become a leader until a new key block is mined. In the meantime, the leader can append microblocks to the chain, which are not subject to PoW mining but are merely hashchained together. As such, microblocks considerably increase the throughput of the whole system and decrease the latency (that said, Bitcoin-NG is still to be stress-tested in practice). In a sense, Bitcoin-NG mixes leader election, often seen in BFT protocols, with a leader-centric protocol in between leader-election epochs. However, what is different in Bitcoin-NG from BFT protocols is that leader election is PoW-based. Consequently, forks are still possible in Bitcoin-NG and consensus finality is not ensured, which may lead to security implications such as asset double-spending, as discussed earlier.

Scaling Blockchain Through Parallelization. Scaling blockchain by making it a blockDAG (directed acyclic graph) rather than a linear chain of blocks, was recently proposed by Lewenberg et al. in the context of PoW [38]. The idea is to allow non-conflicting transactions (e.g., those transactions that do not constitute double-spending attempts) to be initially on different forks, but to eventually merge the forks by mining a block that would include them both in the ledger.⁶ The BFT and state-machine replication communities have also been intensively exploring the idea of *parallel replication* for a few years now, leveraging parallelization of execution of independent requests (transactions) (see, e.g., [32, 42]).

Eliminating Communication and Resource Overhead in BFT Protocols. As we have already discussed, the major challenge for BFT protocols that prevents their wider adoption in blockchain is their scalability in terms of the number of nodes. Stellar [43] is an ongoing effort aimed at removing unanimously accepted membership lists from BFT protocols, while maintaining the other BFT advantages over PoW. Other approaches target the BFT scalability without changing membership assumptions. These include *optimistic BFT* protocols [3, 51] which feature linear communication complexity in the “common case” and resort to expensive $O(n^2)$ communication among nodes featured by classical protocols such as PBFT [10] only if the network and the process fault pattern are particularly unfavorable. However, even optimistic BFT have a resource and communication overhead when compared to crash-tolerant replication protocols (e.g., [30, 36, 49]), which are better proven in practice and may serve as a baseline for BFT.

⁵ <https://etherchain.org/statistics/basic>.

⁶ BlockDAGs are conceptually similar to the notion of parallel sharded chains (sidechains) combined with merge mining.

To rectify this, Liu et al. recently proposed a novel network and node fault model called XFT [39] that allows one to tolerate up to $n/2$ Byzantine nodes. At the same time, XFT features message patterns characteristic to crash-tolerant replication protocols, i.e., without the overhead pertaining to typical BFT message patterns. To this end, XFT (“cross” fault tolerance) challenges the established ability of a BFT adversary to control the network and Byzantine nodes simultaneously, decoupling network faults from Byzantine-node faults, treating them as largely independent. As such, XFT goes in the direction of a more realistic adversary model that resembles the one of PoW blockchains, which are not very concerned with the ability of the adversary to control the entire communication network.

Finally, another appealing direction for future BFT-based blockchain is BFT protocols that leverage small pieces of trusted hardware (e.g., [31]) to improve communication and reduce resource cost.

Randomized BFT. Randomized BFT protocols (e.g., [7,9,55]) are appealing alternative to standard, eventually synchronous [20] BFT protocols such as PBFT. Specifically, randomized BFT protocols circumvent the FLP consensus impossibility result [24] by guaranteeing correctness with very high probability (i.e., always, except with negligible probability), rather than deterministically. This allows randomized BFT protocols to be completely asynchronous [4].

For many years, an issue with randomized BFT protocols has been their performance. Specifically, classical randomized BFT (e.g., [4,7,9,55]) are very inefficient compared to eventually synchronous, deterministic BFT protocols mostly due to overhead of cryptographic tools they use. However, this may be changing soon with novel randomized BFT protocols such as HoneyBadger [46] showing promise for good practical performance (i.e., reasonably high throughput) with up to about 100 nodes, through cherry-picking best available cryptographic tools for randomization as well as processing requests in very large batches. Clearly, large batches negatively impact latency, but this could be addressed by *Hybrid BFT* protocols [2] that may combine very efficient optimistic and deterministic BFT protocols (e.g., those described in [3]) with practical randomized protocols such as HoneyBadger. Early examples of such Hybrid BFT protocols can be found in [2,35,51], but the development of future Hybrid BFT protocols can be facilitated by using the modular BFT design framework described in [3].

Mixing PoW and BFT. Recently, Decker et al. [15] have proposed to enhance PoW blockchain with BFT (concretely, the PBFT protocol [10]), primarily to ensure consensus finality in a PoW blockchain by using BFT. SCP [41] also proposes a hybrid PoW/BFT protocol, using PoW for identity management and (parallel and hierarchical) BFT consensus for agreement. Clearly, the above discussion on the importance of scaling BFT in terms of the number of nodes is also critical to such approaches that mix PoW and BFT.

4 Conclusion and Open Problems

We briefly overviewed state of the art as well as emerging directions towards scalable blockchain. We contrasted proof-of-work (PoW) and Byzantine fault-tolerant (BFT) consensus protocols, highlighting their respective advantages.

Future work will be very dynamic and interesting. Making Fig. 1 more precise, i.e., placing various protocols at the correct place with respect to their performance versus their node-scalability, entails a fair amount of research, but represents an immediate open problem that needs to be better understood to facilitate future blockchain scalability improvements. Furthermore, a lot of potential lies in synergies between PoW and BFT, both when it comes to combining protocol techniques and when it comes to refining the adversarial and network models.

Finally, for the most demanding blockchain applications, it would be interesting to move computationally expensive parts of BFT protocols (e.g., cryptography) closer to hardware. In general, implementing consensus in hardware is indeed very appealing and may yield impressive performance, as attested by recent proposals that explore this idea in the context of crash fault-tolerance [29, 50].

References

1. The “Zeitgeist attack”. <http://bitcoin.stackexchange.com/questions/1055/what-is-the-zeitgeist-attack-does-it-affect-all-blockchain-technologies>
2. Aguilera, M.K., Toueg, S.: Failure detection and randomization: a hybrid approach to solve consensus. *SIAM J. Comput.* **28**(3), 890–903 (1998)
3. Aublin, P.-L., Guerraoui, R., Knežević, N., Quéma, V., Vukolić, M.: The next 700 BFT protocols. *ACM Trans. Comput. Syst.* **32**(4), 12:1–12:45 (2015)
4. Ben-Or, M.: Another advantage of free choice: completely asynchronous agreement protocols (extended abstract). In: *Proceedings of the Second Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC)*, pp. 27–30 (1983)
5. Bessani, A.N., Sousa, J., Alchieri, E.A.P.: State machine replication for the masses with BFT-SMART. In: *44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2014*, pp. 355–362 (2014)
6. Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J.A., Felten, E.W.: Sok: research perspectives and challenges for Bitcoin and cryptocurrencies. In: *2015 IEEE Symposium on Security and Privacy, SP 2015*, pp. 104–121 (2015)
7. Bracha, G.: An asynchronous $[(n-1)/3]$ -resilient consensus protocol. In: *Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pp. 154–162 (1984)
8. Brewer, E.A.: Towards robust distributed systems (abstract). In: *ACM Symposium on Principles of Distributed Computing (PODC)*, p. 7 (2000)
9. Cachin, C., Kursawe, K., Shoup, V.: Random oracles in constantipole: practical asynchronous byzantine agreement using cryptography (extended abstract). In: *Proceedings of the Nineteenth Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pp. 123–132 (2000)
10. Castro, M., Liskov, B.: Practical Byzantine fault tolerance and proactive recovery. *ACM Trans. Comput. Syst.* **20**(4), 398–461 (2002)

11. Chandra, T.D., Griesemer, R., Redstone, J.: Paxos made live: an engineering perspective. In: Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC). ACM (2007)
12. Clement, A., Wong, E., Alvisi, L., Dahlin, M., Marchetti, M.: Making Byzantine fault tolerant systems tolerate Byzantine faults. In: Proceedings of the 6th USENIX symposium on Networked systems design and implementation, NSDI 2009, pp. 153–168. USENIX Association (2009)
13. Corbett, J.C., Dean, J., Epstein, M., et al.: Spanner: google’s globally distributed database. *ACM Trans. Comput. Syst. (TOCS)* **31**(3), 8 (2013)
14. de Prisco, R.: On Building Blocks for Distributed Systems. Ph.D. thesis, Massachusetts Institute of Technology (1999)
15. Decker, C., Seidel, J., Wattenhofer, R.: Bitcoin meets strong consistency. In: 17th International Conference on Distributed Computing and Networking (ICDCN) (2016)
16. Decker, C., Wattenhofer, R.: Information propagation in the Bitcoin network. In: 13th IEEE International Conference on Peer-to-Peer Computing, IEEE P2P 2013, pp. 1–10 (2013)
17. Défago, X., Schiper, A., Urbán, P.: Total order broadcast and multicast algorithms: taxonomy and survey. *ACM Comput. Surv.* **36**(4), 372–421 (2004)
18. Douceur, J.R.: The sybil attack. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, pp. 251–260. Springer, Heidelberg (2002)
19. Jiaqing, D., Sciascia, D., Elnikety, S., Zwaenepoel, W., Pedone, F.: Clock-RSM: low-latency inter-datacenter state machine replication using loosely synchronized physical clocks. In: The 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) (2014)
20. Dwork, C., Lynch, N., Stockmeyer, L.: Consensus in the presence of partial synchrony. *J. ACM* **35**(2), 288–323 (1988)
21. Dwork, C., Naor, M.: Pricing via processing or combatting junk mail. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 139–147. Springer, Heidelberg (1993)
22. Eyal, I., Gencer, A.E., Sirer, E.G., van Renesse, R.: Bitcoin-NG: a scalable blockchain protocol. In: 13th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2016 (2016)
23. Eyal, I., Sirer, E.G.: Majority is not enough: bitcoin mining is vulnerable. In: Christin, N., Safavi-Naini, R. (eds.) FC 2014. LNCS, vol. 8437, pp. 436–454. Springer, Heidelberg (2014)
24. Fischer, M.J., Lynch, N.A., Paterson, M.S.: Impossibility of distributed consensus with one faulty process. *J. ACM* **32**(2), 374–382 (1985)
25. Garay, J., Kiayias, A., Leonardos, N.: The bitcoin backbone protocol: analysis and applications. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 281–310. Springer, Heidelberg (2015)
26. Gervais, A., Karame, G.O., Capkun, V., Capkun, S.: Is Bitcoin a decentralized currency? *IEEE Secur. Priv.* **12**(3), 54–60 (2014)
27. Gilbert, S., Lynch, N.A.: Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News* **33**(2), 51–59 (2002)
28. Guerraoui, R., Vukolić, M.: Refined quorum systems. *Distrib. Comput.* **23**(1), 1–42 (2010)
29. Istvan, Z., Sidler, D., Alonso, G., Vukolić, M.: Consensus in a box: inexpensive coordination in hardware. In: Proceedings of the 13th USENIX Symposium on Networked Systems Design and Implementation, NSDI (2016)

30. Junqueira, F.P., Reed, B.C., Serafini, M.: Zab: high-performance broadcast for primary-backup systems. In: Proceedings of the Conference on Dependable Systems and Networks (DSN), pp. 245–256 (2011)
31. Kapitzka, R., Behl, J., Cachin, C., Distler, T., Kuhnle, S., Mohammadi, S.V., Schröder-Preikschat, W., Stengel, K.: CheapBFT: resource-efficient Byzantine fault tolerance. In: European Conference on Computer Systems, Proceedings of the Seventh EuroSys Conference 2012, pp. 295–308 (2012)
32. Kaprutos, M., Wang, Y., Quema, V., Clement, A., Alvisi, L., Dahlin, M.: All about Eve: execute-verify replication for multi-core servers. In: Proceedings of the 10th USENIX Conference on Operating Systems Design and Implementation, OSDI 2012, pp. 237–250. USENIX Association (2012)
33. Karame, G.O., Androutaki, E., Roeschlin, M., Gervais, A., Capkun, S.: Misbehavior in bitcoin: a study of double-spending and accountability. *ACM Trans. Inf. Syst. Secur.* **18**(1), 2 (2015)
34. Kotla, R., Alvisi, L., Dahlin, M., Clement, A., Wong, E.: Zyzyva: speculative byzantine fault tolerance. *ACM Trans. Comput. Syst.* **27**, 7:1–7:39 (2010)
35. Kursawe, K., Shoup, V.: Optimistic asynchronous atomic broadcast. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 204–215. Springer, Heidelberg (2005)
36. Lamport, L.: The part-time parliament. *ACM Trans. Comput. Syst.* **16**, 133–169 (1998)
37. Lamport, L., Shostak, R., Pease, M.: The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.* **4**, 382–401 (1982)
38. Lewenberg, Y., Sompolinsky, Y., Zohar, A.: Inclusive block chain protocols. In: Böhme, R., Okamoto, T. (eds.) FC 2015. LNCS, vol. 8975, pp. 528–547. Springer, Heidelberg (2015)
39. Liu, S., Cachin, C., Quéma, V., Vukolić, M.: XFT: practical fault tolerance beyond crashes. CoRR, abs/1502.05831 (2015)
40. Losa, G.: Modularity in the design of robust distributed algorithms. Ph.D. thesis, Ecole Polytechnique Federale de Lausanne (2014)
41. Luu, L., Narayanan, V., Baweja, K., Zheng, C., Gilbert, S., Saxena, P.: SCP: a computationally-scalable Byzantine consensus protocol for blockchains. Cryptology ePrint Archive, Report 2015/1168 (2015). <http://eprint.iacr.org/>
42. Marandi, P.J., Bezerra, C.E.B., Pedone, F.: Rethinking state-machine replication for parallelism. In: IEEE 34th International Conference on Distributed Computing Systems, ICDCS 2014, pp. 368–377 (2014)
43. Mazières, D.: The Stellar consensus protocol: A federated model for internetlevelconsensus, November 2015. <https://www.stellar.org/papers/stellar-consensus-protocol.pdf>
44. Mickens, J., The saddest moment. *Login Usenix Mag.* **39**(3) (2014)
45. Milkau, U., Bott, J.: Digitalisation in payments: From interoperability to centralised models? *J. Payment Strategy Syst.* **9**(3), 321–340 (2015)
46. Miller, A., Xia, Y., Croman, K., Shi, E., Song, D.: The honeybadger of BFT protocols. In: Cryptology ePrint Archive 2016/199 (2016)
47. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system, May 2009
48. O’Dwyer, K.J., Malone, D.: Bitcoin mining and its energy footprint. In: Proceedings of the 2014 IET Irish Signals & Systems Conference (2014)
49. Ongaro, D., Ousterhout, J.: In search of an understandable consensus algorithm. In: Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference, USENIX ATC 2014, pp. 305–320. USENIX Association (2014)

50. Poke, M., Hoefler, T.: DARE: high-performance state machine replication on RDMA networks. In: Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing, HPDC 2015, pp. 107–118 (2015)
51. Ramasamy, H.G.V., Cachin, C.: Parsimonious asynchronous byzantine-fault-tolerant atomic broadcast. In: Anderson, J.H., Prencipe, G., Wattenhofer, R. (eds.) OPODIS 2005. LNCS, vol. 3974, pp. 88–102. Springer, Heidelberg (2006)
52. Rodrigues, R., Liskov, B., Chen, K., Liskov, M., Schultz, D.A.: Automatic reconfiguration for large-scale reliable storage systems. *IEEE Trans. Dependable Sec. Comput.* **9**(2), 145–158 (2012)
53. Schneider, F.B.: Implementing fault-tolerant services using the state machine approach: a tutorial. *ACM Comput. Surv.* **22**(4), 299–319 (1990)
54. Sompolinsky, Y., Zohar, A.: Secure high-rate transaction processing in bitcoin. In: Böhme, R., Okamoto, T. (eds.) FC 2015. LNCS, vol. 8975, pp. 507–527. Springer, Heidelberg (2015)
55. Toueg, S.: Randomized byzantine agreements. In: Proceedings of the Third Annual ACM Symposium on Principles of Distributed Computing, Vancouver, B.C., Canada, 27–29 August 1984, pp. 163–178 (1984)
56. Vukolić, M.: Quorum Systems: With Applications to Storage and Consensus. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, San Rafael (2012)
57. Wood, G.: Ethereum: A secure decentralised generalised transaction ledger (2015). <http://gavwood.com/paper.pdf>
58. Zohar, A.: Bitcoin: under the hood. *Commun. ACM* **58**(9), 104–113 (2015)