An In-Depth Look of BFT Consensus in Blockchain: Challenges and Opportunities

Suyash Gupta, Jelle Hellings, Sajjad Rahnama, Mohammad Sadoghi

Exploratory Systems Lab

Department of Computer Science

University of California, Davis

Abstract

Since the introduction of Bitcoin—the first wide-spread application driven by blockchains—the interest of the public and private sector in blockchains has skyrocketed. At the core of this interest are the ways in which blockchains can be used to improve data management, e.g., by enabling federated data management via decentralization, resilience against failure and malicious actors via replication and consensus, and strong data provenance via a secured immutable ledger.

In practice, high-performance blockchains for data management are usually built in permissioned environments in which the participants are vetted and can be identified. In this setting, blockchains are typically powered by Byzantine fault-tolerant consensus protocols. These consensus protocols are used to provide full replication among all honest blockchain participants by enforcing an unique order of processing incoming requests among the participants.

In this tutorial, we take an in-depth look at Byzantine fault-tolerant consensus. First, we take a look at the theory behind replicated computing and consensus. Then, we delve into how common consensus protocols operate. Finally, we take a look at current developments and briefly look at our vision moving forward.

ACM Reference Format:

Suyash Gupta, Jelle Hellings, Sajjad Rahnama, Mohammad Sadoghi. 2019. An In-Depth Look of BFT Consensus in Blockchain: Challenges and Opportunities. In *20th International Middleware Conference Tutorials (Middleware Tutorials '19)*, December 9–13, 2019, Davis, CA, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3366625.3369437>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Middleware Tutorials '19, December 9–13, 2019, Davis, CA, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-7040-0/19/12...\$15.00

<https://doi.org/10.1145/3366625.3369437>

1 Introduction

Since the introduction of *Bitcoin*—the first wide-spread application driven by *blockchain*—the interest of the public and private sector in blockchain has skyrocketed. Recently, this interest has cumulated in the introduction of several blockchain-inspired database systems and blockchain fabrics [2, 3, 25, 26, 52, 53]. Blockchain-based systems have also been demonstrated to address challenges in various other fields such as the trade of valuable commodities [11, 15, 61], food production [29], managing land property rights [59], managing identities [4, 11, 59], supporting transparent aid delivery [4, 59], managing health care data [7, 31, 43], insurance fraud prevention [42], energy production and energy trading [61], and managing compliance with the GDPR [16].

In each of these systems and use cases, blockchain technology is used to improve data management in one way or another. Blockchain technology—which is at the intersection of database systems, distributed systems, and cryptography—provides *promises for new directions in data management* by utilizing and combining underused techniques from distributed systems and cryptography in novel and powerful ways [10, 19, 40, 51].

The role of Bitcoin, Ethereum, and other cryptocurrencies in these developments cannot be understated: at their core, Bitcoin and Ethereum provided proof of the viability of blockchain techniques in non-trivial large-scale settings. At the core of these cryptocurrencies is the maintenance of a *blockchain* among a group of replicas. This blockchain holds an ordered record of all transactions and this record is secured against changes using cryptographic primitives. To assure that all replicas agree on the same set of transactions and maintain the same blockchain, new transactions are agreed upon via a consensus protocol (which are fault-tolerant counterparts of the classical two-phase and three-phase commit protocols used in distributed database systems [32, 37, 63]). From our perspective, these technologies can strengthen data management in three vital directions:

1. The blockchain structure in Bitcoin and Ethereum utilizes cryptographic primitives to harden against unwanted changes (providing immutability). These techniques can also be used to secure a traditional database journal against tampering, which is a major step towards providing systems that provide reliable *data*

provenance. Such tamper-resistant data storage can be used to irrefutably proof any claims about the current state of the data.

2. Fault-tolerant consensus—used to replicate the transactions that are part of the blockchain—is designed to deal with *malicious behavior* of some of the replicas. The same techniques can also be used to deal with malicious behavior caused by the compromise (e.g., due to a cyberattack) of some of the replicas in a system. In this way, blockchain technology promises one way to harden against cyberattacks on data-based services, which can reduce the huge societal and economic impacts of such attacks [33, 54–56, 64].
3. Finally, fault-tolerance consensus can also be used as the technique for supporting *federated data management*, the collective management of a single database among various stakeholders. Federated data management is in itself a major step towards dealing with *data quality issues* arising from the non-federated interchange of information between various stakeholders and, as such, can reduce the huge negative economic impact of bad data [24, 41, 62]. Moreover, in this federated setting, the immutable and irrefutable structure of the blockchain can further help in *policing disputes* in cases where some stakeholders do not trust each other.

The explosion of blockchain based applications, products, and proof-of-concepts (see, e.g., [3, 10, 50, 69]), has led to the development of several different approaches towards consensus. We can roughly categorize these approaches into *permissionless blockchains*, as used in public settings by Bitcoin and other cryptocurrencies, and *permissioned blockchains*, which are better suited for managed private environments.

In *permissionless blockchains* such as Bitcoin, *Proof-of-Work*-inspired consensus algorithms are used to replicate data [36, 50, 69]. These algorithms require *limited communication* between replicas and can operate in unstructured peer-to-peer networks in which independent parties can join and leave at any time [60]. Proof-of-Work uses computationally complex puzzles to limit the influence any malicious party has on the evolution of the blockchain. At the same time, these puzzles incur a high *computational costs* on all parties, which has raised questions about the sustainability of the *energy consumption* of permissionless blockchain systems [17, 68]. Additionally, the complexity of Proof-of-Work puzzles causes relative long transaction processing times (minutes to hours) and significantly limits the number of transactions a permissionless blockchain can handle: in 2017, it was reported that Bitcoin can only process 7 transactions per second, whereas Visa already processes 2000 transactions per second on average [59]. Finally, the design of Proof-of-Work prevents scalability, as adding more computational power to the network will only increase the cost

of Proof-of-Work puzzles. Other permissionless consensus algorithms based on Proof-of-Work, such as Proof-of-Space and Proof-of-Stake have similar limitations with respect to resource usage and transaction throughput.

Most data management use cases do not require the flexibility of unstructured peer-to-peer networks in which participants can join and leave at any time: data management systems are usually employed in a managed environment in which all participants are known, can be identified, and are vetted. This is exactly the setting for which *permissioned blockchains* are designed. In these blockchains, traditional Byzantine fault-tolerant replication techniques based on *consensus algorithms* such as PBFT are employed to accept, order, and execute client transactions among all replicas [5, 6, 13, 45, 46, 57, 66, 67]. The benefit of these consensus algorithms, compared to Proof-of-Work-based algorithms, is that they have low computational costs, low transaction processing times, and high transaction throughput. This is already exemplified in 2002 by *BFS*, a fault-tolerant version of the networked file system [38], which could already handle hundreds of transactions per second [13, 14].

In this tutorial, we will provide a deep dive into consensus protocols with a focus on data management. To do so, we take an in-depth look at Byzantine fault-tolerant consensus protocols, the main technique powering permissioned blockchains.

Our tutorial will focus on three avenues. First, we look at the theoretical framework in which permissioned blockchains operate. Then, we look at practical high-performance consensus protocols and at current developments. Finally, we look at the challenges in the design of future-proof high-performance permissioned blockchain systems that can deal with huge amounts of data, and provide our vision on future developments. These avenues are detailed in the following three sections. Finally, in Section 5, we provide a summary of the topics addressed during the tutorial.

2 Fault-tolerant distributed computing

Blockchains are, at their basis, fully replicated distributed systems that aim to maintain data consistency. The well-known CAP Theorem puts restrictions on the types of failures these blockchains can deal with while guaranteeing continued services [8, 9, 30]. The CAP Theorem puts rather general limitations on the design of blockchains, however. More specific limitations are also known, as the Byzantine consensus problem and other related problems, such as the Byzantine agreement problem and the interactive consistency problem, have received considerable attention.

It is well-known that the Byzantine agreement problem can only be solved when using synchronous communication [28, 49, 65]. In a synchronous environment with n replicas of which f are Byzantine (e.g., malicious), Byzantine

agreement requires that $n > 3f$ [20, 21]. When strong cryptographic primitives are available, this can be improved to $n > f$ [23, 47, 58] (although practical systems will still require $n > 2f$). Additionally, bounds on the amount of communication and the quality of the network are known [18, 20–23, 27].

3 Practical consensus protocols

Having provided a theoretical background, we make the step toward detailing practical consensus protocols. We do so by a full coverage of the *Practical Byzantine Fault Tolerant* consensus protocol (PBFT) of Castro et al. [12–14]. Next, we also look at the lineage of consensus protocols that refine and improve PBFT. This detailed overview will cover many of the practical consensus protocols currently in use and, simultaneously, also covers recent developments. Our coverage will include protocols such as PBFT, HotStuff [70], Zyzzyva [5, 44, 45], FaB [48], SynBFT [1], RBFT [6], PoE, and MultiBFT [34].

4 Challenges and our Vision

As outlined above, the approaches taken by permissionless and permissioned blockchains towards fault-tolerant replication have *benefits* and practical use cases. Unfortunately, fault-tolerant replication is *challenged* by the scalability and performance required by many modern big-data-driven applications. In specific, we see that there is no obvious way to scale up fault-tolerant replication: adding more replicas will only increase the cost of replication and decrease the throughput of the system, even when using the most efficient consensus protocols. We will close our tutorial by discussing recent steps toward the design of new fault-tolerant architectures that step away from the full-replicated nature of blockchains, this to increase scalability and the ability to serve big-data-driven applications. In specifics, we will look at two low-level techniques, *cluster-sending* [39] and *delayed-replication*, and at a high-performance sharded and geo-scale aware architecture that is enabled by these techniques.

5 Outline of the Tutorial

The tutorial starts with a general-purpose introduction to blockchains from the perspective of data management. Next, each of three avenues discussed in Section 2–4 will receive attention. In specific, we plan to explore the following outline:

Blockchains and Data Management

1. What are blockchain and why should I care about them.
2. Blockchains from the perspective of data management.
3. The connection between blockchains and consensus.
4. Permissioned and permissionless blockchain systems.

Overview of Resilient Distributed Systems

1. What are distributed and resilient systems.
2. General limitations of distributed systems: the CAP Theorem.
3. Failure models: crashes versus Byzantine failures.
4. Communication models: asynchronous versus synchronous communication.
5. Consensus, Byzantine agreement, and interactive consistency.
6. Theoretical results on resilient distributed systems.

Byzantine Fault-Tolerance in Practice

1. The Practical Byzantine Fault Tolerant consensus algorithm in detail.
2. Improving on PBFT: digital signatures versus symmetric encryption, multi-round reasoning, optimistic execution, randomization, threshold signatures, trusted components, parallelization, and other techniques.

Challenges and our Vision

1. Towards sharding and geo-scale designs.
2. Specialized designs via cluster-sending.
3. Specialized designs via Byzantine learners.

This tutorial is based on the outline of our upcoming book on fault-tolerant transaction processing on blockchains [35].

References

- [1] Ittai Abraham, Srinivas Devadas, Danny Dolev, Kartik Nayak, and Ling Ren. 2018. Synchronous Byzantine Agreement with Expected $O(1)$ Rounds, Expected $O(n^2)$ Communication, and Optimal Resilience. (2018).
- [2] Mohammad Javad Amiri, Divyakant Agrawal, and Amr El Abbadi. 2019. CAPER: A Cross-application Permissioned Blockchain. *Proceedings of the VLDB Endowment* 12, 11 (2019), 1385–1398. <https://doi.org/10.14778/3342263.3342275>
- [3] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muradharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolić, Sharon Weed Cocco, and Jason Yellick. 2018. Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. In *Proceedings of the Thirteenth EuroSys Conference*. ACM, 30:1–30:15. <https://doi.org/10.1145/3190508.3190538>
- [4] GSM Association. 2017. Blockchain for Development: Emerging Opportunities for Mobile, Identity and Aid. (2017). <https://www.gsma.com/mobilefordevelopment/wp-content/uploads/2017/12/Blockchain-for-Development.pdf>
- [5] Pierre-Louis Aublin, Rachid Guerraoui, Nikola Knežević, Vivien Quéma, and Marko Vukolić. 2015. The Next 700 BFT Protocols. *ACM Transactions on Computer Systems* 32, 4 (2015), 12:1–12:45. <https://doi.org/10.1145/2658994>
- [6] Pierre-Louis Aublin, Sonia Ben Mokhtar, and Vivien Quéma. 2013. RBFT: Redundant Byzantine Fault Tolerance. In *2013 IEEE 33rd International Conference on Distributed Computing Systems*. IEEE, 297–306. <https://doi.org/10.1109/ICDCS.2013.53>

- [7] Burkhard Blechschmidt. 2018. *Blockchain in Europe: Closing the Strategy Gap*. Technical Report. Cognizant Consulting. <https://www.cognizant.com/whitepapers/blockchain-in-europe-closing-the-strategy-gap-codex3320.pdf>
- [8] Eric Brewer. 2012. CAP twelve years later: How the “rules” have changed. *Computer* 45, 2 (2012), 23–29. <https://doi.org/10.1109/MC.2012.37>
- [9] Eric A. Brewer. 2000. Towards Robust Distributed Systems (Abstract). In *Proceedings of the Nineteenth Annual ACM Symposium on Principles of Distributed Computing*. ACM, 7–7. <https://doi.org/10.1145/343477.343502>
- [10] Christian Cachin and Marko Vukolic. 2017. Blockchain Consensus Protocols in the Wild (Keynote Talk). In *31st International Symposium on Distributed Computing (Leibniz International Proceedings in Informatics)*, Vol. 91. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 1:1–1:16. <https://doi.org/10.4230/LIPIcs.DISC.2017.1>
- [11] Michael Casey, Jonah Crane, Gary Gensler, Simon Johnson, and Neha Narula. 2018. *The Impact of Blockchain Technology on Finance: A Catalyst for Change*. Technical Report. International Center for Monetary and Banking Studies. https://www.cimb.ch/uploads/1/1/5/4/115414161/geneva21_1.pdf
- [12] Miguel Castro. 2001. *Practical Byzantine Fault Tolerance*. Ph.D. Dissertation. Massachusetts Institute of Technology. <http://hdl.handle.net/1721.1/86581>
- [13] Miguel Castro and Barbara Liskov. 1999. Practical Byzantine Fault Tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*. USENIX Association, 173–186.
- [14] Miguel Castro and Barbara Liskov. 2002. Practical Byzantine Fault Tolerance and Proactive Recovery. *ACM Transactions on Computer Systems* 20, 4 (2002), 398–461. <https://doi.org/10.1145/571637.571640>
- [15] Christie’s. 2018. Major Collection of the Fall Auction Season to be Recorded with Blockchain Technology. (2018). https://www.christies.com/presscenter/pdf/9160/RELEASE_ChristiesxArtoxyxEbworth_9160_1.pdf
- [16] Cindy Compert, Maurizio Luinetti, and Bertrand Portier. 2018. *Blockchain and GDPR: How blockchain could address five areas associated with GDPR compliance*. Technical Report. IBM Security. <https://public.dhe.ibm.com/common/ssi/ecm/61/en/61014461usen/security-ibm-security-solutions-wg-white-paper-external-61014461usen-20180319.pdf>
- [17] Alex de Vries. 2018. Bitcoin’s Growing Energy Problem. *Joule* 2, 5 (2018), 801–805. <https://doi.org/10.1016/j.joule.2018.04.016>
- [18] Richard A. DeMillo, Nancy A. Lynch, and Michael J. Merritt. 1982. Cryptographic Protocols. In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*. ACM, 383–400. <https://doi.org/10.1145/800070.802214>
- [19] Tien Tuan Anh Dinh, Rui Liu, Meihui Zhang, Gang Chen, Beng Chin Ooi, and Ji Wang. 2018. Untangling Blockchain: A Data Processing View of Blockchain Systems. *IEEE Transactions on Knowledge and Data Engineering* 30, 7 (2018), 1366–1385. <https://doi.org/10.1109/TKDE.2017.2781227>
- [20] D. Dolev. 1981. Unanimity in an unknown and unreliable environment. In *22nd Annual Symposium on Foundations of Computer Science*. IEEE, 159–168. <https://doi.org/10.1109/SFCS.1981.53>
- [21] Danny Dolev. 1982. The Byzantine generals strike again. *Journal of Algorithms* 3, 1 (1982), 14–30. [https://doi.org/10.1016/0196-6774\(82\)90004-9](https://doi.org/10.1016/0196-6774(82)90004-9)
- [22] Danny Dolev and Rüdiger Reischuk. 1985. Bounds on Information Exchange for Byzantine Agreement. *J. ACM* 32, 1 (1985), 191–204. <https://doi.org/10.1145/2455.214112>
- [23] D. Dolev and H. Strong. 1983. Authenticated Algorithms for Byzantine Agreement. *SIAM J. Comput.* 12, 4 (1983), 656–666. <https://doi.org/10.1137/0212045>
- [24] Wayne W. Eckerson. 2002. *Data quality and the bottom line: Achieving Business Success through a Commitment to High Quality Data*. Technical Report. The Data Warehousing Institute, 101communications LLC.
- [25] Muhammad El-Hindi, Carsten Binnig, Arvind Arasu, Donald Kossmann, and Ravi Ramamurthy. 2019. BlockchainDB: A Shared Database on Blockchains. *Proceedings of the VLDB Endowment* 12, 11 (2019), 1597–1609. <https://doi.org/10.14778/3342263.3342636>
- [26] Muhammad El-Hindi, Martin Heyden, Carsten Binnig, Ravi Ramamurthy, Arvind Arasu, and Donald Kossmann. 2019. BlockchainDB – Towards a Shared Database on Blockchains. In *Proceedings of the 2019 International Conference on Management of Data*. ACM, 1905–1908. <https://doi.org/10.1145/3299869.3320237>
- [27] Michael J. Fischer and Nancy A. Lynch. 1982. A lower bound for the time to assure interactive consistency. *Inform. Process. Lett.* 14, 4 (1982), 183–186. [https://doi.org/10.1016/0020-0190\(82\)90033-3](https://doi.org/10.1016/0020-0190(82)90033-3)
- [28] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. 1985. Impossibility of Distributed Consensus with One Faulty Process. *J. ACM* 32, 2 (1985), 374–382. <https://doi.org/10.1145/3149.214121>
- [29] Lan Ge, Christopher Brewster, Jacco Spek, Anton Smeenk, and Jan Top. 2017. *Blockchain for agriculture and food: Findings from the pilot study*. Technical Report. Wageningen University. <https://www.wur.nl/nl/Publicatie-details.htm?publicationId=publication-way-353330323634>
- [30] Seth Gilbert and Nancy Lynch. 2002. Brewer’s Conjecture and the Feasibility of Consistent, Available, Partition-tolerant Web Services. *SIGACT News* 33, 2 (2002), 51–59. <https://doi.org/10.1145/564585.564601>
- [31] William J. Gordon and Christian Catalini. 2018. Blockchain Technology for Healthcare: Facilitating the Transition to Patient-Driven Interoperability. *Computational and Structural Biotechnology Journal* 16 (2018), 224–230. <https://doi.org/10.1016/j.csbj.2018.06.003>
- [32] Jim Gray. 1978. Notes on Data Base Operating Systems. In *Operating Systems, An Advanced Course*. Springer-Verlag, 393–481. https://doi.org/10.1007/3-540-08755-9_9
- [33] Andy Greenberg. 2018. The Untold Story of NotPetya, the Most Devastating Cyberattack in History. (2018). <https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/>
- [34] Suyash Gupta, Jelle Hellings, and Mohammad Sadoghi. 2019. Brief Announcement: Revisiting Consensus Protocols through Wait-Free Parallelization. In *33rd International Symposium on Distributed Computing (DISC 2019) (Leibniz International Proceedings in Informatics (LIPIcs))*, Vol. 146. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 44:1–44:3. <https://doi.org/10.4230/LIPIcs.DISC.2019.44>
- [35] Suyash Gupta, Jelle Hellings, and Mohammad Sadoghi. 2020. *Fault-Tolerant Distributed Transactions on Blockchains*. (to appear).
- [36] Suyash Gupta and Mohammad Sadoghi. 2018. *Blockchain Transaction Processing*. Springer International Publishing, 1–11. https://doi.org/10.1007/978-3-319-63962-8_333-1
- [37] Suyash Gupta and Mohammad Sadoghi. 2018. EasyCommit: A Non-blocking Two-phase Commit Protocol. In *Proceedings of the 21st International Conference on Extending Database Technology*. Open Proceedings, 157–168. <https://doi.org/10.5441/002/edbt.2018.15>
- [38] Thomas Haynes and David Noveck. 2015. RFC 7530: Network File System (NFS) Version 4 Protocol. (2015). <https://tools.ietf.org/html/rfc7530>
- [39] Jelle Hellings and Mohammad Sadoghi. 2019. Brief Announcement: The Fault-Tolerant Cluster-Sending Problem. In *33rd International Symposium on Distributed Computing (DISC 2019) (Leibniz International Proceedings in Informatics (LIPIcs))*, Vol. 146. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 45:1–45:3. <https://doi.org/10.4230/LIPIcs.DISC.2019.45>
- [40] Maurice Herlihy. 2019. Blockchains from a Distributed Computing Perspective. *Commun. ACM* 62, 2 (2019), 78–85. <https://doi.org/10.1145/3209623>
- [41] Thomas N. Herzog, Fritz J. Scheuren, and William E. Winkler. 2007. *Data Quality and Record Linkage Techniques*. Springer New York. <https://doi.org/10.1007/0-387-69505-2>

- [42] Matt Higginson, Johannes-Tobias Lorenz, Björn Mäijstermann, and Peter Braad Olesen. 2017. *The promise of blockchain*. Technical Report. McKinsey&Company. <https://www.mckinsey.com/industries/financial-services/our-insights/the-promise-of-blockchain>
- [43] Maged N. Kamel Boulos, James T. Wilson, and Kevin A. Clauson. 2018. Geospatial blockchain: promises, challenges, and scenarios in health and healthcare. *International Journal of Health Geographics* 17, 1 (2018), 1211–1220. <https://doi.org/10.1186/s12942-018-0144-x>
- [44] Ramakrishna Kotla, Lorenzo Alvisi, Mike Dahlin, Allen Clement, and Edmund Wong. 2007. Zyzzyva: Speculative Byzantine Fault Tolerance. In *Proceedings of Twenty-first ACM SIGOPS Symposium on Operating Systems Principles*. ACM, 45–58. <https://doi.org/10.1145/1294261.1294267>
- [45] Ramakrishna Kotla, Lorenzo Alvisi, Mike Dahlin, Allen Clement, and Edmund Wong. 2009. Zyzzyva: Speculative Byzantine Fault Tolerance. *ACM Transactions on Computer Systems* 27, 4 (2009), 7:1–7:39. <https://doi.org/10.1145/1658357.1658358>
- [46] Leslie Lamport. 1978. The implementation of reliable distributed multiprocess systems. *Computer Networks (1976)* 2, 2 (1978), 95–114. [https://doi.org/10.1016/0376-5075\(78\)90045-4](https://doi.org/10.1016/0376-5075(78)90045-4)
- [47] Leslie Lamport, Robert Shostak, and Marshall Pease. 1982. The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems* 4, 3 (1982), 382–401. <https://doi.org/10.1145/357172.357176>
- [48] Jean-Philippe Martin and Lorenzo Alvisi. 2006. Fast Byzantine Consensus. *IEEE Transactions on Dependable and Secure Computing* 3, 3 (2006), 202–215. <https://doi.org/10.1109/TDSC.2006.35>
- [49] Shlomo Moran and Yaron Wolfstahl. 1987. Extended impossibility results for asynchronous complete networks. *Inform. Process. Lett.* 26, 3 (1987), 145–151. [https://doi.org/10.1016/0020-0190\(87\)90052-4](https://doi.org/10.1016/0020-0190(87)90052-4)
- [50] Satoshi Nakamoto. 2009. Bitcoin: A Peer-to-Peer Electronic Cash System. (2009). <https://bitcoin.org/bitcoin.pdf>
- [51] Arvind Narayanan and Jeremy Clark. 2017. Bitcoin’s Academic Pedigree. *Commun. ACM* 60, 12 (2017), 36–45. <https://doi.org/10.1145/3132259>
- [52] Senthil Nathan, Chander Govindarajan, Adarsh Saraf, Manish Sethi, and Praveen Jayachandran. 2019. Blockchain Meets Database: Design and Implementation of a Blockchain Relational Database. *Proceedings of the VLDB Endowment* 12, 11 (2019), 1539–1552. <https://doi.org/10.14778/3342263.3342632>
- [53] Faisal Nawab and Mohammad Sadoghi. 2019. Blockplane: A Global-Scale Byzantizing Middleware. In *35th International Conference on Data Engineering (ICDE)*. IEEE, 124–135. <https://doi.org/10.1109/ICDE.2019.00020>
- [54] Dick O’Brien. 2017. *Internet Security Threat Report: Ransomware 2017, An ISTR Special Report*. Technical Report. Symantec. <https://www.symantec.com/content/dam/symantec/docs/security-center/white-papers/istr-ransomware-2017-en.pdf>
- [55] The Council of Economic Advisers. 2018. *The Cost of Malicious Cyber Activity to the U.S. Economy*. Technical Report. Executive Office of the President of the United States. <https://www.whitehouse.gov/wp-content/uploads/2018/03/The-Cost-of-Malicious-Cyber-Activity-to-the-U.S.-Economy.pdf>
- [56] National Audit Office. 2018. Investigation: WannaCry cyber attack and the NHS. (2018). <https://www.nao.org.uk/report/investigation-wannacry-cyber-attack-and-the-nhs/>
- [57] Michael Okun. 2016. *Byzantine Agreement*. Springer New York, 255–259. https://doi.org/10.1007/978-1-4939-2864-4_60
- [58] M. Pease, R. Shostak, and L. Lamport. 1980. Reaching Agreement in the Presence of Faults. *J. ACM* 27, 2 (1980), 228–234. <https://doi.org/10.1145/322186.322188>
- [59] Michael Pisa and Matt Juden. 2017. *Blockchain and Economic Development: Hype vs. Reality*. Technical Report. Center for Global Development. <https://www.cgdev.org/publication/blockchain-and-economic-development-hype-vs-reality>
- [60] Bitcoin Project. 2018. Bitcoin Developer Guide: P2P Network. (2018). <https://bitcoin.org/en/developer-guide#p2p-network>
- [61] PwC. 2016. Blockchain – an opportunity for energy producers and consumers? (2016). <https://www.pwc.com/gx/en/industries/energy-utilities-resources/publications/opportunity-for-energy-producers.html>
- [62] Thomas C. Redman. 1998. The Impact of Poor Data Quality on the Typical Enterprise. *Commun. ACM* 41, 2 (1998), 79–82. <https://doi.org/10.1145/269012.269025>
- [63] Dale Skeen. 1982. *A Quorum-Based Commit Protocol*. Technical Report. Cornell University.
- [64] Symantec. 2018. Internet Security Threat Report, Volume 32. (2018). <https://www.symantec.com/content/dam/symantec/docs/reports/istr-23-2018-en.pdf>
- [65] Gadi Taubenfeld and Shlomo Moran. 1996. Possibility and impossibility results in a shared memory environment. *Acta Informatica* 33, 1 (1996), 1–20. <https://doi.org/10.1007/s002360050034>
- [66] Gerard Tel. 2001. *Introduction to Distributed Algorithms* (2nd ed.). Cambridge University Press.
- [67] Maarten van Steen and Andrew S. Tanenbaum. 2017. *Distributed Systems* (3th ed.). Maarten van Steen. <https://www.distributed-systems.net/>
- [68] Harald Vranken. 2017. Sustainability of bitcoin and blockchains. *Current Opinion in Environmental Sustainability* 28 (2017), 1–9. <https://doi.org/10.1016/j.cosust.2017.04.011>
- [69] Gavin Wood. 2016. Ethereum: a secure decentralised generalised transaction ledger. (2016). <https://gavwood.com/paper.pdf> EIP-150 revision.
- [70] Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan Gueta, and Ittai Abraham. 2019. HotStuff: BFT Consensus with Linearity and Responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. ACM, 347–356. <https://doi.org/10.1145/3293611.3331591>