

# **High Performance, Low Energy, and Trustworthy Blockchains Using Satellites**

**Other titles in Foundations and Trends® in Networking**

*Contagion Source Detection in Epidemic and Infodemic Outbreaks:  
Mathematical Analysis and Network Algorithms*

Chee Wei Tan and Pei-Duo Yu

ISBN: 978-1-63828-250-1

*Distributed Coding in A Multiple Access Environment*

Yanru Tang, Faeze Heydaryan and Jie Luo

ISBN: 978-1-68083-468-0

*Age of Information: A New Concept, Metric, and Tool*

Antzela Kosta, Nikolaos Pappas and Vangelis Angelakis

ISBN: 978-1-68083-360-7

*Network and Protocol Architectures for Future Satellite Systems*

Tomaso de Cola, Alberto Ginesi, Giovanni Giambene,

George C. Polyzos, Vasilios A. Siris, Nikos Fotiou and Yiannis Thomas

ISBN: 978-1-68083-334-8

*Duality of the Max-Plus and Min-Plus Network Calculus*

Jorg Liebeherr

ISBN: 978-1-68083-294-5

# High Performance, Low Energy, and Trustworthy Blockchains Using Satellites

---

**Dennis Shasha**

New York University

**Taegyun Kim**

New York University

**Joseph Bonneau**

New York University

**Yan Michalevsky**

Cryptosat Inc.

**Gil Shotan**

Stanford University

**Yonatan Winetraub**

Cryptosat Inc.

**now**

the essence of knowledge

Boston — Delft

## Foundations and Trends<sup>®</sup> in Networking

*Published, sold and distributed by:*

now Publishers Inc.  
PO Box 1024  
Hanover, MA 02339  
United States  
Tel. +1-781-985-4510  
[www.nowpublishers.com](http://www.nowpublishers.com)  
[sales@nowpublishers.com](mailto:sales@nowpublishers.com)

*Outside North America:*

now Publishers Inc.  
PO Box 179  
2600 AD Delft  
The Netherlands  
Tel. +31-6-51115274

The preferred citation for this publication is

D. Shasha *et al.*. *High Performance, Low Energy, and Trustworthy Blockchains Using Satellites*. Foundations and Trends<sup>®</sup> in Networking, vol. 13, no. 4, pp. 252–325, 2023.

ISBN: 978-1-63828-275-4

© 2023 D. Shasha *et al.*

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: [www.copyright.com](http://www.copyright.com)

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; [www.nowpublishers.com](http://www.nowpublishers.com); [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, [www.nowpublishers.com](http://www.nowpublishers.com); e-mail: [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

# **Foundations and Trends<sup>®</sup> in Networking**

Volume 13, Issue 4, 2023

## **Editorial Board**

### **Editor-in-Chief**

**Sanjay Shakkottai**

The University of Texas at Austin

United States

## Editorial Scope

### Topics

Foundations and Trends® in Networking publishes survey and tutorial articles in the following topics:

- Modeling and Analysis of:
  - Ad Hoc Wireless Networks
  - Sensor Networks
  - Optical Networks
  - Local Area Networks
  - Satellite and Hybrid Networks
  - Cellular Networks
  - Internet and Web Services
- Protocols and Cross-Layer Design
- Network Coding
- Energy-Efficiency Incentives/Pricing/Utility-based
- Games (co-operative or not)
- Security
- Scalability
- Topology
- Control/Graph-theoretic models
- Dynamics and Asymptotic Behavior of Networks

### Information for Librarians

Foundations and Trends® in Networking, 2023, Volume 13, 4 issues. ISSN paper version 1554-057X. ISSN online version 1554-0588. Also available as a combined paper and online subscription.

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Our Contribution: The Bounce Protocol . . . . .	6
1.2	System Components . . . . .	7
1.3	What Are the Advantages of Satellites over Terrestrial Blockchain Arbiters? . . . . .	9
1.4	High Level Security Comparison of Bounce With Other Protocols . . . . .	10
<b>2</b>	<b>Bounce Blockchain with Never-Traitorous Satellites</b>	<b>12</b>
2.1	Strawman Multi-Satellite Protocol . . . . .	13
2.2	Trust Model for the Full Protocol . . . . .	16
2.3	Data Structures . . . . .	18
2.4	State of the Blockchain . . . . .	23
2.5	Slot Protocol . . . . .	24
2.6	Handling Failures . . . . .	28
2.7	Reset Protocol . . . . .	29
2.8	Satellite Role in Reset Process . . . . .	31
2.9	Statechart Representation of the Satellite Protocol . . . . .	33
2.10	Broadcast Ground Stations . . . . .	34
2.11	Statechart Representation of the Mission Control Protocols . . . . .	38
2.12	Guarantees . . . . .	41

2.13 Mission Control Reliability . . . . .	43
2.14 Denial-of-Service Attacks on Mission Control . . . . .	45
<b>3 Liftoff, Mission Control, Orbit Design, and Slot Timing</b>	<b>46</b>
3.1 Liftoff . . . . .	46
3.2 Periodic Mission Control . . . . .	46
3.3 Orbit Design and Satellite Timing . . . . .	48
3.4 Sizes of Messages . . . . .	51
3.5 Slot Length . . . . .	53
<b>4 Alternative Blockchain Designs</b>	<b>56</b>
4.1 Terrestrial Systems to Support Blockchains . . . . .	56
<b>5 Handling Traitorous Bounce Units</b>	<b>60</b>
5.1 Failures to Avoid . . . . .	61
5.2 Flock Protocol . . . . .	62
5.3 State Diagram for Bounce Units . . . . .	65
<b>6 Conclusion</b>	<b>67</b>
<b>Acknowledgements</b>	<b>69</b>
<b>References</b>	<b>70</b>



# High Performance, Low Energy, and Trustworthy Blockchains Using Satellites

Dennis Shasha<sup>1</sup>, Taegyun Kim<sup>1</sup>, Joseph Bonneau<sup>1</sup>, Yan Michalevsky<sup>2</sup>, Gil Shotan<sup>3</sup> and Yonatan Winetraub<sup>2</sup>

<sup>1</sup>*Department of Computer Science, New York University, USA; shasha@courant.nyu.edu, taegyun.kim@nyu.edu, jcb@cs.nyu.edu*

<sup>2</sup>*Cryptosat Inc., USA; yan@cryptosat.io, yonatan@cryptosat.io*

<sup>3</sup>*Stanford University, USA; gilsho@cs.stanford.edu*

---

## ABSTRACT

Blockchains are meant to provide an append-only sequence (ledger) of transactions. Security commonly relies on a consensus protocol in which forks in the sequence are either prevented completely or are exponentially unlikely to last more than a few blocks. This monograph proposes the design of algorithms and a system to achieve high performance (a few seconds from the time of initiation for transactions to enter the blockchain), the absence of forks, and a very low energy cost (a per transaction cost that is a factor of a billion or more less than bitcoin).

The foundational component of this setup is a group of satellites whose blockchain protocol code can be verified and burned into read-only memory. Because such satellites can perhaps be destroyed but cannot be captured (unlike even fortified terrestrial servers), a reasonable assumption is that

---

Dennis Shasha, Taegyun Kim, Joseph Bonneau, Yan Michalevsky, Gil Shotan and Yonatan Winetraub (2023), “High Performance, Low Energy, and Trustworthy Blockchains Using Satellites”, *Foundations and Trends® in Networking*: Vol. 13, No. 4, pp 252–325. DOI: 10.1561/13000000070.

©2023 D. Shasha *et al.*

the blockchain protocol code in the satellites may fail to make progress either permanently or intermittently but will not be traitorous.

A second component of this setup is a group of terrestrial sites whose job is to broadcast information about blocks and to summarize the blockchain ledger. These can be individuals who are eager to get a fee for service. Even if many of these behave traitorously (against their interests as fee-collectors), a small number of honest ones is sufficient to ensure safety and liveness.

A third component of this setup is a Mission Control entity which will act very occasionally to assign roles to terrestrial sites and time slots to satellites. These assignments will be multisigned using the digital signatures of a widely distributed group of human governors. A reasonable assumption on Mission Control is that, for reputational reasons, they will not send any signed message that would either contradict a previous message or attest to an incorrect affirmation. Because Mission Control needs to act very infrequently (to a first approximation, only when satellites fail), any actions of Mission Control can be carefully and publicly scrutinized.

Given these components and these reasonable assumptions, our protocol, called **Bounce**, will achieve ledger functionality for arbitrarily sized blocks at under five seconds per block (based on experiments done with the International Space Station) and at negligible energy cost.

This monograph will discuss the overall architecture and algorithms of such a system, the assumptions it makes, and the guarantees it gives.

Keywords: *blockchain, satellites, security, proof-of-work, proof-of-stake*

---

# 1

---

## Introduction

---

The essential functionality of a blockchain is to provide a single, immutable history with a total ordering over all transactions that are committed. We will refer to this immutable total ordering as the **blockchain ledger** or **ledger** for short.

Here the word “single” (or “unforked”) means that if principal  $A$  sees a block of transactions  $b_1$  before  $b_2$  in the history (or “blockchain”) then any other principal  $B$  will see that same relationship in the history. Some blockchains, notably Bitcoin and others relying on Nakamoto-style consensus, provide a slightly weaker “common prefix” property [25] guaranteeing only that all parties will agree to a long, growing prefix of the chain with high probability. The word “immutable” (also called “append-only”) means that once  $b_1$  and  $b_2$  are in the history, they will remain in the history and their order will remain the same as well.

Introducing forks into a blockchain can be profitable for malicious actors. To give an emotive example, Bob could purchase gems from Alice, register that purchase on a block, leave with the gems, and then cause a fork in the blockchain with the result that the longest chain in that blockchain would not contain the block recording that purchase. In such a case, Bob could spend the same funds a second time (and even more times). This is called a **double-spending attack**.

Because the ability to introduce long forks into a blockchain is potentially both very profitable for malicious actors and very disruptive to normal users of the blockchain, any blockchain protocol must either prevent forking or mask it. To achieve this, blockchains invoke consensus protocols that control how blocks are added to the blockchain.

The protocols are divided into two categories:

- “permissionless” (or “decentralized”) protocols in which any principal can potentially add blocks to the blockchain; or
- “permissioned” protocols in which only certain pre-selected principals are allowed to do so.

### **Permissionless (Decentralized) Approaches**

Among the many proposed permissionless approaches [5], [24], perhaps the two most popular are:

1. Proof-of-work, in which the solution to a computationally difficult problem determines the right to add a block on the blockchain. This is the core idea behind the original Bitcoin protocol [30] and many popular successors such as the former Ethereum [39] protocol. This approach can fail if malicious principals control most of the computing power. It also incurs significant recurring energy costs (as of this writing, 127 terawatt-hours, which is more than Norway) which scale at least loosely in proportion to the value protected by the system [6], [9].
2. Proof-of-stake, in which principals reach consensus on new blocks using a voting mechanism with votes proportional to monetary stake in the system. This approach can allow forks if holders of a majority (or sometimes just 1/3) of the stake conspire to act maliciously. Forking can be made even easier if there are honest-but-inactive principals who fail to participate.

In addition, there are many subtle attacks specific to proof-of-stake protocols, including nothing-at-stake attacks [7], retroactive (“long-range”) stake compromise [17], stake-bleeding [26], race-to-the-door takeovers [6], and selfish endorsing [31]. The incentive

analysis is also quite complex and depends intricately on the reward schedule [35], token valuation [22] and returns to lending [13]. Ethereum moved to proof of stake in September 2022 [10], [21]. Ethereum is so large that it will be difficult for coalitions of delegates to gain control. Smaller stake-based systems are much more vulnerable to such an attack. Further, as even Ethereum undergoes continuous development, software changes could make new attacks possible.

Proof-of-stake can also fail due to long-range attacks in which a malicious party gathers secret key material from former participants in the system who no longer have any incentive to behave honestly [4], [17].

### Permissioned (Centralized or Semi-centralized) Approaches

1. Centralized approaches, the opposite extreme from fully permissionless consensus, trust a single central entity typically called a *bank* [15]. This is the simplest and most efficient approach possible and was implicit in early digital cash proposals in the 1980s and 1990s, but requires strong assumptions about correctness of the central bank. Furthermore, the bank is able to arbitrarily censor certain transaction types from the network, whereas censorship-resistance is often stated as a goal of permissionless cryptocurrency systems. For these reasons there are few serious proposals for cryptocurrencies with a single central bank.
2. Semi-centralized approaches, in which consensus is determined by a supermajority of well-known and semi-trusted principals rather than a single one. Most prominent are the Hyperledger Framework [23] and the Facebook Diem project (formerly Libra) [1]. Hyperledger offers a variety of consensus mechanisms: (i) Kafka in which a leader orders transactions (ii) Redundant Byzantine Fault Tolerance [2] (iii) and trusted computing environment-based systems (e.g. based on Intel's SGX). Diem proposes using roughly a dozen semi-trusted nodes running a new Byzantine fault-tolerant consensus protocol, HotStuff, designed for high speed in

the blockchain use case [40]. The Avalanche protocol[32] supports a directed acyclic graph rather than a chain. Its fundamental building block, called the Snowball protocol, can be used to support a chain. Its safety depends on at least partial synchrony and the absence of partitions. All of these proposals require the careful selection of the semi-trusted entities, as a malicious coalition of critical size can fork the system.

## 1.1 Our Contribution: The Bounce Protocol

All of the above consensus approaches can fail in certain circumstances. Bounce uses satellites to create a hard-to-touch chain arbiter that can ensure consensus by itself (in a modified semi-centralized way) or can be added as an extra layer to any consensus protocol to prevent forks.

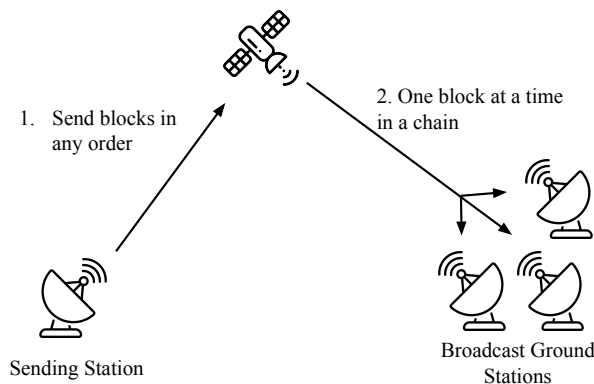
In its simplest idealized form, Bounce consists of a single satellite which receives blocks, signs them and returns them (“bounces” them back) to earth, ensuring that only a single block is signed for each time slot. See Figure 1.1.

In that setting, the satellite is a centralized server with the modification that it can be designed to have its (very simple) protocol software burned into a Read-Only Memory to be both public and unmodifiable once in space. In that way, its ability to *form a total order among transactions* can easily be verified. We call this scheme **Publicly Verifiable Centralized Consensus**. Note that while we envision that the Bounce protocol code will be in read-only memory, the flight control code for the navigational computer may be modifiable.

While the single satellite would maintain a chain without forks, it is unrealistic for two reasons:

1. It assumes the satellite never fails.
2. Unless the satellite is geosynchronous over a well-equipped area of the earth or there are relay satellites, there will be gaps in its coverage (e.g. when the satellite passes over the ocean).

Thus, our actual proposal incorporates multiple low-earth orbiting satellites to provide primarily for fault tolerance but also (even in the



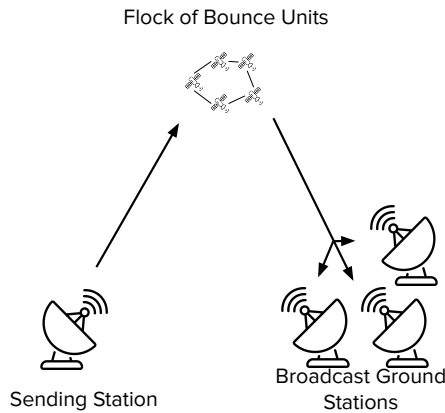
**Figure 1.1:** Idealized Bounce Blockchain consists of a perfectly reliable and trustworthy satellite to which any Sending Station on earth could send a block, the satellite would order the blocks, and send them back to one or more Broadcast Ground Stations which broadcast the numbered and signed blocks widely. The technical goal of this monograph is to achieve this same functionality with failure-prone satellites and other devices.

absence of a relay architecture) for the sake of continuous or near-continuous coverage. The net result will be a low latency blockchain without forks that supports ledger functionality. Bounce does this at negligible energy cost.

## 1.2 System Components

Concretely, the continuously active components of the real Bounce Blockchain system (as opposed to the strawman one) consist of (i) several satellites in low earth orbit where each satellite carries one or more cryptographic units (called Bounce Units) each having the functionality of a hardware security module in the spirit of [29]; (ii) a set of terrestrial Sending Stations that receive transactions from end users, package them into blocks and relay (digests of) blocks to the Bounce Units (iii) a set of terrestrial Broadcast Ground Stations that receive messages from the Bounce Units and from Sending Stations; and (iv) a communications infrastructure of the Broadcast Ground Stations supporting both communication among these principals and sending information broadly to users of the Blockchain. See Figure 1.2.

In addition, there is a Mission Control component which assigns roles to the various principals.



**Figure 1.2:** Components of the Bounce Blockchain system: (i) Satellites (typically small ones like Cubesats) containing Bounce Units to sign blocks. (ii) Sending Stations to upload transactions to Bounce Units. (iii) Broadcast Ground Stations to receive messages from Bounce Units. (iv) A communications infrastructure among Ground Stations. (vi) Mission Control (not shown) which assigns roles to various principals and which assigns time slots to various Bounce Units.

We make some assumptions about the components that come either from underlying technology or the publicly verifiable governance of this system:

1. All honest Principals know the public keys of the Mission Control servers.
2. There is a public key infrastructure (PKI), managed by Mission Control, which ensures that each “Principal” (Bounce Unit, Sending Station, and Broadcast Ground Station) has a public key known to all other Principals. Note that the PKI need not extend to end users of the Bounce ecosystem.
3. The Mission Control Administrators will determine for each Principal  $A$  whether  $A$  is a Sending Station, Broadcast Ground Station, or Bounce Unit. Some principals may take on several roles. Each Principal must know the public keys of the other Principals.



### 1.3. *What Are the Advantages of Satellites over Terrestrial Blockchain Arbiters?*

9

4. When there are many satellites, Mission Control will determine which satellite (and which Bounce Units on that satellite if there are several) is responsible for which time slot, where each time slot corresponds to an interval when a block will be added to the blockchain. All Principals know these assignments.

Section 3.1 discusses the pragmatics of achieving these assumptions.

### 1.3 **What Are the Advantages of Satellites over Terrestrial Blockchain Arbiters?**

A satellite architecture for blockchains yields several benefits:

- Naturally broadcast. A satellite can send to many places on earth in one messages without depending on fibre or other wiring. This reduces the possibility of isolating it. For example, low earth orbiting satellites at only 600 kilometers have a coverage of several million square kilometers [11], giving them many possibilities to communicate.
- A disadvantage of satellites is that their signals can be attenuated due to rain, fog, and snow. That is why the protocol uses digests of transaction blocks instead of blocks. This requires very modest bandwidths (under a kilobit per second during normal operation) as discussed in Section 3.4.
- Satellites are difficult to destroy and even more difficult to capture physically given today's technology, removing the need for guards. State actors can destroy satellites, but they would have a hard time gaining physical possession and any such attempt would be extremely obvious. By contrast, it is feasible to gain access to a terrestrial data center, even a guarded one. Both terrestrial data centers and satellites can make the protocol processors tamper-resistant.

Note however that the protocols we describe here could apply to terrestrial counterparts that would play the part of satellites. Such terrestrial counterparts would not require launching, so would be cheaper

to establish. They would require guards though, so we think satellites would ultimately be less expensive.

#### 1.4 High Level Security Comparison of Bounce With Other Protocols

Because we haven't yet given the details of the Bounce Protocol, we will simply affirm some properties. You the reader can see that they hold as you read the monograph.

- Basically all protocols have a set of administrators who have to be trusted to maintain software, set policy, etc. The reason this is tolerated is that the intervention of those administrators is “punctuated” – that is it occurs only rarely so in principle is open to the public eye. In practice this may be difficult because software changes can have subtle effects.
- To avoid forks, proof of work systems make no assumptions about the honesty of individuals, but do make the assumption that no group of malicious agents control a majority of the computing power [25], [30]. Such a group of colluding malicious agents could cause forks.
- To avoid forks, proof of stake systems assume that no group of malicious agents control a majority of the stake [14], [27]. Sometimes the assumption is stronger, viz. the malicious agents are assumed to control less than 1/3 of the stake [8], [10].
- The Bounce system described in this monograph assumes arbiters (which we conceive of as satellites for the reasons offered in the previous section) programmed with a protocol that can be burned in using a Read-only memory. That protocol can be publicly verified. Bounce also assumes the existence of a group of responsible individuals whom we term collectively as Mission Control. The Mission Control members correspond to system administrators. The Mission Control Administrators multisign their messages in a distributed fashion. Further Mission Control is rarely active

(punctuated): only when a satellite fails or many Sending Stations fail, so its actions (which consist of role assignments) can be easily monitored. Unlike software updates which can be subtle even to experts, the Mission Control messages are easy to understand by interested (even inexpert) parties. Because their bad actions can be so easily discovered, the Mission Control administrators will be reluctant to cheat. Provided Mission Control behaves properly and the minimal assumptions on the satellites hold, there will be no forks.

In summary, the Bounce approach achieves the energy efficiency of proof of stake with more effective security.

## References

---

- [1] Z. Amsden, R. Arora, S. Bano, M. Baudet, S. Blackshear, A. Bothra, G. Cabrera, C. Catalini, K. Chalkias, E. Cheng, A. Ching, A. Chursin, G. Danezis, G. D. Giacomo, D. L. Dill, H. Ding, N. Doudchenko, V. Gao, Z. Gao, F. Garillot, M. Gorven, P. Hayes, J. M. Hou, Y. Hu, K. Hurley, K. Lewi, C. Li, Z. Li, D. Malkhi, S. Margulis, B. Maurer, P. Mohassel, L. d. Naurois, V. Nikolaenko, T. Nowacki, O. Orlov, D. Perelman, A. Pott, B. Proctor, S. Qadeer, Rain, D. Russi, B. Schwab, S. Sezer, A. Sonnino, H. Venter, L. Wei, N. Wernerfelt, B. Williams, Q. Wu, X. Yan, T. Zakian, and R. Zhou, *The Libra Blockchain*, 2020. URL: <https://diem-developers-components.netlify.app/papers/the-diem-blockchain/2020-05-26.pdf> (accessed on 03/23/2023).
- [2] P. Aublin, S. Mokhtar, and V. Quéma, “Rbft: Redundant byzantine fault tolerance,” 2013, pp. 297–306. URL: <http://dx.doi.org/10.1109/ICDCS.2013.53>.
- [3] Y. Aumann and Y. Lindell, “Security against covert adversaries: Efficient protocols for realistic adversaries,” *Journal of Cryptology*, vol. 23, no. 2, 2010, pp. 281–343.
- [4] S. Azouvi, G. Danezis, and V. Nikolaenko, “Winkle: Foiling long-range attacks in proof-of-stake systems,” in *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pp. 189–201, 2020.

- [5] S. Bano, A. Sonnino, M. Al-Bassam, S. Azouvi, P. McCorry, S. Meiklejohn, and G. Danezis, “Sok: Consensus in the age of blockchains,” in *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, pp. 183–198, 2019.
- [6] J. Bonneau, “Hostile blockchain takeovers (short paper),” in *International Conference on Financial Cryptography and Data Security*, Springer, pp. 92–100, 2018.
- [7] 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*, pp. 459–473, 2019.
- [8] E. Buchman, J. Kwon, and Z. Milosevic, “The latest gossip on BFT consensus,” *CoRR*, vol. abs/1807.04938, 2018. URL: <http://arxiv.org/abs/1807.04938>.
- [9] E. Budish, “The economic limits of bitcoin and the blockchain,” National Bureau of Economic Research, Tech. Rep., 2018.
- [10] V. Buterin, D. Hernandez, T. Kamphofner, K. Pham, Z. Qiao, D. Ryan, J. Sin, Y. Wang, and Y. X. Zhang, “Combining GHOST and casper,” *CoRR*, vol. abs/2003.03052, 2020. URL: <https://arxiv.org/abs/2003.03052>.
- [11] S. Cakaj, B. Kamo, A. Lala, and A. Rakipi, “The coverage analysis for low earth orbiting satellites at low elevation,” *International Journal of Advanced Computer Science and Applications*, vol. 5, 2014. DOI: [10.14569/IJACSA.2014.050602](https://doi.org/10.14569/IJACSA.2014.050602).
- [12] M. Castro and B. Liskov, “Practical byzantine fault tolerance,” in *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, ser. OSDI '99, pp. 173–186, New Orleans, Louisiana, USA: USENIX Association, 1999. URL: <http://dl.acm.org/citation.cfm?id=296806.296824>.
- [13] T. Chitra, “Competitive equilibria between staking and on-chain lending,” *arXiv preprint arXiv:2001.00919*, 2019.

- [14] P. Daian, R. Pass, and E. Shi, “Snow white: Robustly reconfigurable consensus and applications to provably secure proof of stake,” in *Financial Cryptography and Data Security - 23rd International Conference*, ser. Lecture Notes in Computer Science, vol. 11598, pp. 23–41, Springer, 2019. DOI: [10.1007/978-3-030-32101-7\\_2](https://doi.org/10.1007/978-3-030-32101-7_2).
- [15] G. Danezis and S. Meikeljohn, “Centrally banked cryptocurrencies,” in *NDSS '16, 21-24 February 2016, San Diego, CA, USA*, 2016. URL: <https://arxiv.org/pdf/1505.06895.pdf>.
- [16] T. Dasu, Y. Kanza, and D. Srivastava, “Unchain your blockchain,” in *Proc. Symposium on Foundations and Applications of Blockchain*, 2018.
- [17] E. Deirmentzoglou, G. Papakyriakopoulos, and C. Patsakis, “A survey on long-range attacks for proof of stake protocols,” *IEEE Access*, vol. 7, 2019, pp. 28 712–28 725.
- [18] A. Delignat-Lavaud, M. Abadi, A. Birrell, I. Mironov, T. Wobber, and Y. Xie, “Web pki: Closing the gap between guidelines and practices.,” in *NDSS*, Citeseer, 2014.
- [19] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The second-generation onion router,” Naval Research Lab Washington DC, Tech. Rep., 2004.
- [20] Ethereum Foundation, *The history of ethereum*, 2021. URL: <https://ethereum.org/en/history/>.
- [21] Ethereum Foundation, *The Merge*, 2023. URL: <https://ethereum.org/en/upgrades/merge/> (accessed on 03/23/2023).
- [22] G. Fanti, L. Kogan, and P. Viswanath, “Economics of proof-of-stake payment systems,” in *Working paper*, 2019.
- [23] L. Foundation, “Hyperledger architecture, volume 1: Introduction to hyperledger business blockchain design philosophy and consensus,” 2017, pp. 1–15. URL: [https://www.hyperledger.org/wp-content/uploads/2017/08/Hyperledger\\_Arch\\_WG\\_Paper\\_1\\_Consensus.pdf](https://www.hyperledger.org/wp-content/uploads/2017/08/Hyperledger_Arch_WG_Paper_1_Consensus.pdf).
- [24] J. Garay and A. Kiayias, “Sok: A consensus taxonomy in the blockchain era,” in *Cryptographers’ Track at the RSA Conference*, Springer, pp. 284–318, 2020.

- [25] J. Garay, A. Kiayias, and N. Leonardos, “The bitcoin backbone protocol: Analysis and applications,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Springer, pp. 281–310, 2015.
- [26] P. Gaži, A. Kiayias, and A. Russell, “Stake-bleeding attacks on proof-of-stake blockchains,” in *2018 Crypto Valley conference on Blockchain technology (CVCBT)*, IEEE, pp. 85–92, 2018.
- [27] A. Kiayias, A. Russell, B. David, and R. Oliynykov, “Ouroboros: A provably secure proof-of-stake blockchain protocol,” in *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference*, ser. Lecture Notes in Computer Science, vol. 10401, pp. 357–388, Springer, 2017. DOI: [10.1007/978-3-319-63688-7\\_12](https://doi.org/10.1007/978-3-319-63688-7_12).
- [28] X. Ling, Z. Gao, Y. Le, L. You, J. Wang, Z. Ding, and X. Gao, “Satellite-aided consensus protocol for scalable blockchains,” *Sensors*, vol. 20, no. 19, 2020. DOI: [10.3390/s20195616](https://doi.org/10.3390/s20195616).
- [29] Y. Michalevsky and Y. Winetraub, “Spacetee: Secure and tamper-proof computing in space using cubesats,” 2017, pp. 1–6. URL: <https://arxiv.org/pdf/1710.01430.pdf>.
- [30] S. Nakamoto, *Bitcoin: A peer-to-peer electronic cash system*, 2008. URL: <https://bitcoin.org/bitcoin.pdf>.
- [31] M. Neuder, D. J. Moroz, R. Rao, and D. C. Parkes, “Selfish behavior in the tezos proof-of-stake protocol,” *arXiv preprint arXiv:1912.02954*, 2019.
- [32] T. Rocket, M. Yin, K. Sekniqi, R. van Renesse, and E. G. Sirer, “Scalable and probabilistic leaderless BFT consensus through metastability,” 2020, pp. 1–21. URL: <https://docs.avax.network/learn/platform-overview/avalanche-consensus>.
- [33] N. Roubini, “Exploring the cryptocurrency and blockchain ecosystem,” *Testimony for the Hearing of the US Senate Committee on Banking, Housing and Community Affairs, October 11, 2018*, 2018, pp. 1–37.
- [34] M. D. Ryan, “Enhanced certificate transparency and end-to-end encrypted mail,” *Cryptology ePrint Archive*, 2013.
- [35] F. Saleh, “Blockchain without waste: Proof-of-stake,” *The Review of financial studies*, vol. 34, no. 3, 2021, pp. 1156–1190.

- [36] B. Schneier, “Using intel’s sgx to attack itself,” 2017. URL: [https://www.schneier.com/blog/archives/2017/03/using%5C\\_intels\\_sg.html](https://www.schneier.com/blog/archives/2017/03/using%5C_intels_sg.html).
- [37] B. Schneier, “Speculation attack against intel’s sgx,” 2018. URL: [https://www.schneier.com/blog/archives/2018/08/speculation\\_att.html](https://www.schneier.com/blog/archives/2018/08/speculation_att.html).
- [38] SpaceChain Foundation, “Spacechain white paper,” 2018, pp. 1–23. URL: <https://spacechain.com/wp-content/uploads/2018/09/SpaceChain-White-Paper.pdf>.
- [39] G. Wood *et al.*, “Ethereum: A secure decentralised generalised transaction ledger,” *Ethereum project yellow paper*, vol. 151, no. 2014, 2014, pp. 1–32.
- [40] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, “Hotstuff: Bft consensus with linearity and responsiveness,” in *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pp. 347–356, 2019.